

CASSINI SPACECRAFT ATTITUDE CONTROL SYSTEM: FLIGHT PERFORMANCE AND LESSONS LEARNED, 1997–2017*

Allan Y. Lee¹ and Thomas A. Burk²

A sophisticated interplanetary spacecraft, Cassini/Huygens was launched on October 15, 1997. Since achieving orbit at Saturn in 2004, Cassini has collected science data throughout its four-year prime mission (2004–08), and has since been approved for first and second extended missions through September 2017. The Cassini Attitude and Articulation Control Subsystem (AACS) is perhaps the spacecraft subsystem that must satisfy the most mission and science pointing requirements. Since launch, the performance of the Cassini AACS design has been superb. All key mission and science requirements are met with significant margins. An overview of the flight performance of the Cassini attitude control system as well as AACS mission operation-centric lessons learned, from launch to 2017, are described by topics. Many of these lessons learned should be applicable to the safe operations of other interplanetary missions. Processes taken by the AACS operation team to guard against “human” errors are also outlined in this paper.

CASSINI/HUYGENS MISSION TO SATURN AND TITAN

The Cassini spacecraft was launched on 15 October 1997 by a Titan 4B launch vehicle. After an interplanetary cruise of 6.7 years, it arrived at Saturn on June 30, 2004. To save propellant, Cassini made several gravity-assist flybys: two at Venus and one each at Earth and Jupiter. Since achieving orbit at Saturn in 2004, Cassini has collected science data throughout its four-year prime mission (2004–08), and has since been approved for first and second extended missions through September 2017. Major science objectives of the Cassini mission include investigations of the configuration and dynamics of Saturn’s magnetosphere, the structure and composition of the rings, the characterization of several of Saturn’s icy satellites, and Titan’s atmosphere constituent abundance.¹

Cassini carries twelve scientific instruments. Six instruments measure properties of objects remote from the spacecraft. An example is the imaging science subsystem with a wide and a narrow angle camera. Cassini also carries six instruments that observe fields, particles, and plasma waves. An example of these instruments is the Ion and Neutral Mass Spectrometer (INMS). INMS is used to determine the chemical, elemental, and isotopic composition of the gaseous and volatile components of the neutral particles and the low-energy ions in Titan’s atmosphere and ionosphere. In 2005, Cassini completed three flybys of Enceladus, a small icy satellite of Saturn. Observations made during these flybys confirmed the existence of watery geysers in the South polar region of Enceladus. The discovery of these watery geysers is an important and unexpected discovery made by Cassini. The Enceladus watery plume is one of the key science investigations of the Cassini Equinox mission (an extension of the Cassini Prime mission, from July 2008 to September 2010).

*Copyright 2017 California Institute of Technology. Government sponsorship acknowledgement.

¹Section Staff, Guidance and Control Section, Division of Autonomous Systems, Jet Propulsion Laboratory, California Institute of Technology, California Institute of Technology, Mail Stop 230-104, 4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA. (818)-354-4097, Allan.Y.Lee@jpl.nasa.gov.

²Cassini Attitude and Articulation Control Subsystem Lead Engineer, Guidance and Control Section, Division of Autonomous Systems, Jet Propulsion Laboratory, California Institute of Technology, California Institute of Technology, Mail Stop 230-104, 4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA. (818)-393-0974, Thomas.A.Burk@jpl.nasa.gov.

In April 2017, the Cassini spacecraft will begin a daring set of ballistic orbits that will hop the rings and dive between the upper atmosphere of Saturn and its innermost D-ring twenty-two times.² During these “proximal” orbits, the spacecraft will be as close as 1,840 km from the 1-bar constant pressure ellipsoidal surface of Saturn. As Cassini plunges past Saturn, the spacecraft will collect rich and valuable data that will better our understanding of Saturn’s gravity and magnetic fields, and atmospheric sciences. The Cassini mission will end on September 15, 2017.

ATTITUDE AND ARTICULATION CONTROL (AACS) SUBSYSTEM

Perhaps no other spacecraft subsystem must satisfy as many science and mission requirements as AACS.¹ The Cassini AACS estimates and controls the attitude of the three-axis stabilized spacecraft. It responds to ground-commanded pointing goals for the spacecraft’s science instruments and communication antennas with respect to targets of interest. To this end, the AACS uses either thrusters or reaction wheels to slew and control the spacecraft attitude. The AACS also executes ground-commanded spacecraft velocity changes. To this end, AACS uses either a gimballed rocket engine or a set of thrusters to make a velocity change. Table 1 lists the key accuracy requirements for the Cassini spacecraft.

Table 1. Key Cassini Accuracy Requirements¹

Pointing Accuracy	Requirements
High-gain antenna (HGA) pointing control requirement (radial 99%)	
X-band (Telecommunications)	3.2 mrad/s
Ka-band (Radio science)	2.0 mrad/s
Ku-band (Radar mapping of Titan)	4.6 mrad/s
S-band (Huygens probe relay tracking)	6.0 mrad/s
Science inertial pointing requirements (radial 99%)	
Control	2.0 mrad/s
Knowledge	1.0 mrad/s
Pointing stability requirements (2σ per axis):	
0.5 s, 1 s	4, 8 μ rad/s
5 s, 22 s, others	36, 100 μ rad/s, others
Gimballed engine ΔV burns (1σ per axis):	
Fixed Magnitude, Fixed Pointing	10 mm/s, 17.5 mm/s
Proportional Magnitude, Proportional Pointing	0.35%, 10 mrad/s
Thruster ΔV burns (1σ per axis):	
Fixed Magnitude, Fixed Pointing	3.5 mm/s, 3.5 mm/s
Proportional Magnitude, Proportional Pointing	2%, 12 mrad/s

To achieve a high degree of maneuverability, and to facilitate high-resolution imaging, Cassini is designed as a three-axis stabilized spacecraft. Attitude determination sensors include two stellar reference units (SRU, prime and backup, star tracker),^{3,4} two Sun sensors (prime and backup),⁵ two Inertial Reference Unit (IRU, each IRU has four 1-dof gyroscopes),⁶ and a 1-dof accelerometer (ACC). Attitude control actuators include four Reaction Wheel Assemblies (RWA, three prime and a backup RWA mounted on an articulate-able platform), two 445-N bi-propellant rocket engines (prime and backup), two sets of 2-dof engine gimbal actuators (EGA, prime and backup), and two sets of 1-N monopropellant thrusters (blow-down, prime and backup).^{7,8} The locations of the Cassini AACS equipment are depicted in Figure 1.¹ The mass of Cassini at launch was 5574 kg of which 3000 kg was main engine bi-propellant consisting of mono-methyl-hydrazine for fuel and nitrogen tetroxide for oxidizer. A single tank with 132.2 kg of hydrazine feeds the thrusters.

AACS acquires stellar reference by first locating the Sun,⁴ and then Sun-pointing the HGA using Sun position knowledge from SSA. The front-end of the Cassini attitude estimator is a pre-filter that combines multiple star updates into one “composite” star update. These composite star updates are sent to the

attitude estimator (an extended Kalman-Bucy filter) every 1-5 seconds. In between star updates, the S/C's attitude is propagated using IRU data. Once attitude is initialized, the AACS maintains knowledge of the spacecraft attitude in a "J2000" celestial coordinate frame. It is defined by the Earth Mean Equator and Equinox at the year 2000 epoch.




Figure 1. Locations of Cassini Attitude Control Equipment¹

During Tour, Saturn, its rings, and its satellites sometime enter the star tracker's FOV. The presence of one or more of these bright (or extended) objects might affect the nominal operation of the star identification (SID) algorithm.⁹ As such, the SID algorithm has to be temporarily "suspended" via an AACS command. While SID is suspended, the S/C attitude is propagated using only data from the gyroscopes. At the start and end times of a "suspend" event, the spacecraft must be quiescent with all per-axis rates <0.01 °/s. The requirement at the start time is imposed to ascertain that the attitude estimator uncertainty is as small as possible for subsequent attitude propagation. The requirement at the end time is imposed to ascertain that the SID algorithm can seamlessly transition from a "suspend" state to the "track" state, without going through any "acquisition" phase.

During early Cruise, Cassini used a set of eight 1-N mono-propellant RCS (reaction control) thrusters to maintain the spacecraft attitude about all axes (see Figure 8). Thrusters were also used to perform the following functions: detumble the spacecraft after it was released from the launch vehicle; control the spacecraft's attitude during low-altitude Titan flybys; bias the RWA's angular momenta; and execute small ΔV burns. A conventional Bang-Off-Bang (BOB) thruster control algorithm is used by Cassini AACS. It uses error signals that are the weighted sums of per-axis attitude errors and attitude rate errors to control thruster firings.¹⁰ However, such a control algorithm can result in "two-sided" limit cycles that waste both hydrazine and thruster on/off cycle. To counter these drawbacks, the Cassini's BOB incorporated a "self-learning" feature to produce, as much as possible, "one-sided" limit cycles in the presence of small environmental torque. The resultant "one-sided" limit cycles save both hydrazine and thruster on/off cycles. See also the section entitled "Tuning of RCS Attitude Controller Parameters."

The Reaction Wheel Assemblies (RWAs) are used primarily for attitude control when precise and stable pointing of a science instrument (such as the narrow angle camera) is required during the prime mission phase. Because the spacecraft's principle axes are very closely aligned with the spacecraft's mechanical axes, the basic structure of the Cassini RWA controller is a decoupled, three-axis, proportional and derivative controller.¹¹ The RWA control torque vector also includes the feed-forward acceleration torque and the gyroscopic torque. Due to the presence of RWA bearing drag torque, a RWA controller with

a “PD” control architecture will not be able to drive the spacecraft attitude control error to zero unless an integral term is added to the controller. This is overcome by the addition of a proportional and integral estimator of the RWA frictional torque. In effect, integral control action is added “locally” to remove any steady-state spacecraft’s attitude control errors. See also the section entitled “RWA Controller Design and Bearing Drag Torque Estimator.”

During the Tour phase, RWA’s must be used to achieve a high level of spacecraft pointing stability needed for imaging operations of cameras and instruments. The use of reaction wheels is subjected to three constraints. Firstly, the spin rates of all RWA must not exceed the angular momentum capacity of the wheels. Secondly, the total number of revolutions of the RWA’s that is incurred as a result of science slews must be kept as low as possible. These two constraints both discourage high-speed wheel operations. Thirdly, the operational hours the wheels spend inside a “low-rpm” region must be minimized. Below a certain spin rate, the thickness of the lubrication film (between the bearing balls and the races) will be smaller than the root-mean-square value of the surface roughness of the balls and races. This will lead to metal-to-metal contact between the balls and the races which is highly undesirable. To operate the RWAs within these constraints, the Cassini AACS team developed and used a ground software tool named Reaction Wheel Bias Optimization Tool (RBOT).¹² See also the section entitled “Reaction Wheel Bias Optimization Tool”.

The Cassini interplanetary mission requires both large and small trajectory correction maneuvers for navigation purposes. Trajectory corrections performed before Saturn Orbit Insertion (SOI) are called Trajectory Correction Maneuvers (TCMs). Corrections performed after SOI are called Orbit Trim Maneuvers (OTMs). Large burns were performed by one of the two rocket engines (with a nominal thrust of 445 N). Smaller ΔV ’s were performed using four Z-facing thrusters (each with a 1-N thrust).^{13–16}

During the Cruise phase, thrusters were used to roll and yaw the spacecraft attitude so as to align the pre-aimed rocket engine with the target ΔV vector. “Settling” times on the order of 5 minutes were “inserted” in between the roll and yaw turns, and in between the end of the yaw turn and the start of the burn. Once a burn was completed, the spacecraft would “un-yaw” and “un-roll” back to its initial attitude. Again, settling times were inserted between these “unwind” turns. These thruster-based slewing imparted unwanted ΔV on the spacecraft. Even though the magnitudes of these ΔV could be predicted, they still, in a small way, affected the accuracy of the burn. As such, beginning with TCM-18 (in April 2002), both the roll and un-roll turns were executed using a set of reaction wheels.

Three-axis stabilized orbiters such as the Viking Mars Orbiter had used two-axis engine gimbal actuator autopilot for Thrust Vector Control (TVC). Cassini uses a similar design. As was practiced on Viking, pre-aim of the gimbal through the predicted spacecraft c.m. was commanded to minimize attitude disturbances at engine ignition. During the burn, the X and Y-axis of the spacecraft’s attitude were controlled by the engine gimbal actuators using a TVC algorithm. At the same time, four Y-facing thrusters were used to control the spacecraft’s Z-axis attitude. The ΔV imparted on the spacecraft is measured by an accelerometer. Since the ACC’s bias changes slightly from burn to burn, it is calibrated before the start of each and every main engine burn. The flight software (FSW) uses the calibrated ACC to accurately determine the magnitude of the ΔV imparted on the spacecraft. The burn is terminated once the commanded ΔV is achieved. Details on the design of the TVC algorithm are given in Reference 17. Flight performance of the TVC algorithm has been excellent and is given in Refs.14–15. See also the section entitled “Engine-based ΔV Control and Operational Issues”.

Cassini uses four Z-facing thrusters to execute small ΔV burns (they are called RCS burns). During a RCS burn, the Z-facing thrusters are used to achieve the targeted ΔV . Engine off-pulsing will be needed to negate the self-induced disturbance torque caused by S/C’s c.m. offset from the Z-axis. In addition, misaligned thrusters will generate disturbance torque about the S/C’s Z-axis. Four Y-facing thrusters are used to control the spacecraft’s Z-axis attitude. See also the section entitled “Thruster-based ΔV Control and Operational Issues”.

Pointing the Cassini spacecraft involves commands to turn to and track targets of interest. Pointing commands do not reference underlying entities like right ascension, declination, or spacecraft body rates. On Cassini, pointing commands reference celestial objects themselves. If the target is Saturn, a single command causes the spacecraft to turn to Saturn. Once there, the spacecraft tracks Saturn until commanded to turn to another object. This means that target motion compensation is inherent in the design. And with more than a dozen science instrument boresights, Cassini pointing commands reference the name of the

science instrument boresight itself as the object to point at Saturn. Cassini is able to explicitly point boresights at celestial objects by use of an innovative pointing model called Inertial Vector Propagation (IVP).¹⁸ IVP provides a remarkably robust yet structured process that takes advantage of powerful onboard vector propagation algorithms combined with sophisticated targeting options that science teams can take full advantage of on the ground.¹⁹

Celestial body objects are represented by inertial vectors which express the actual position in the solar system of Cassini with respect to Saturn, for example, or Earth with respect to the Sun. These inertial vectors can thus range from unit vectors to vectors many millions of kilometers in magnitude. Each inertial vector “head” has a time-varying position and velocity with respect to its vector “base”. Inertial vectors are used to establish a pointing direction by simple vector addition. For example, to point at the Earth while in Saturn orbit, the following is used (see Figure 2):




Figure 2. Pointing Cassini at Earth Using IVP Vector Addition

The three vectors in the above addition are individually propagated within the IVP flight software and are considered “active” for as long as Cassini-to-Earth is the desired pointing direction. Active vectors are part of a pointing “tree” of inertial vectors and require onboard propagation throughout their period of activity. The Sun is the root of the inertial vector tree.

Pointing the Cassini spacecraft is achieved using two “inertial” vectors and two “body” vectors to command a unique inertial attitude in space. For example, the X-Band radio-frequency boresight vector (XBAND) of the HGA is defined as a unit vector in the spacecraft body coordinate frame. XBAND is an example of a primary body vector. The HGA itself points parallel to the -Z spacecraft body vector, while the XBAND is the electrical boresight and is slightly offset from -Z. The fundamental requirement is to point a selected primary body vector at a primary inertial vector object. For example, during data playback to Earth, the primary pointing is commanded to be “XBAND” to “EARTH” where the Earth is the head of an inertial vector from the spacecraft to Earth. To establish a complete 3-axis inertial attitude a secondary pair of body and inertial vectors are specified. For example, a secondary pair could be “NAC” to “SATURN” where NAC is the narrow angle camera boresight body vector, and Saturn is the time-varying inertial vector from the spacecraft to the center of Saturn. The primary pointing is always achieved, while the angle between the secondary pair is minimized given the constraint that the primary vectors must be collinear. In practice, this means the secondary body vector is commanded into the plane defined by the primary and secondary inertial vectors (see Figure 3). This defines the “twist” around the primary body vector.

Inertial vectors reside in an onboard inertial vector table that is populated by ground IVP commands. A ground IVP tool is loaded with the desired pointing sequence and with up-to-date spacecraft and celestial body ephemeris and physical constants files. The tool “fits” inertial vectors to approximate the ephemeris using conic or Chebyshev polynomial coefficients. Conic coefficients are just position, velocity, and a gravity parameter. Conic propagation involves solving Kepler’s equation.^{47–48} Conics propagate accurately when there is a dominant central body (for Cassini, the dominant body is Saturn). Because conics do not always suffice, Chebyshev polynomial vectors (of up to 12th order) are also used by IVP.⁴⁹ Polynomials are

used during transitional periods in the spacecraft's trajectory as gravitational dominance shifts from one body to another.

Conic and polynomial vector fits are done on the ground using the IVP tool. Vectors active for long periods of time require multiple "segments" where each segment is fit to 40 μ rad maximum error where the error is defined as the angular difference between the propagation and the ephemeris data. Each vector segment is a separate IVP command with a start and end time. The IVP tool defines a new segment before the prior segment has "expired" (reached its end time). Body vectors can also be updated, for example after an in-flight calibration. All IVP vectors are merged with the background sequence prior to uplink to Cassini.




Figure 3. Pointing Body Vectors at Inertial Vectors

For Cassini AACS, two aspects of pointing commands must be checked by the so-called Constraint Monitor (CMT) before they are implemented by the attitude controller.²⁰ First, the per-axis slew rates and accelerations are checked against a set of rate and acceleration limits that represents the capability of the thrusters (or a set of reaction wheels). Any per-axis rates or accelerations that is higher than its CMT limits will be "truncated" by the CMT. Only the truncated slew command is used by the attitude controller. CMT also checks to make sure that the angle between a specific body vector and an inertial vector (usually the Sun-line vector) is larger than a pre-selected threshold. For example, the bore-sight vector of the narrow angle camera must NOT be closer than 12° relative to the Sun-line vector (actually, any part of the Sun). The commanded S/C attitude will be altered by the CMT if a geometric constraint violation is anticipated by the CMT. See also the section entitled "Ground Simulation and Tools – Keys to Finding Errors".

Pre-launch, many engineers worked on the designs and testing of AACS hardware and software. Knowledge of the limitations of the attitude control subsystem must be recorded and passed to the mission operations team to ensure spacecraft safety.²¹ One efficient way to transfer the knowledge is for the design and test teams to write "flight rules" for the Operations team to follow during mission operations. The rationales behind the needs to enforce these rules must be clearly stated. Flight rules must be enforced until they are no longer applicable. Approved flight rules could be enforced either via ground software, manual inspection, or by other means. A flight rule is waived only if there is a strong motivation to do that and when the consequences of violation are fully understood. For example, launch-related flight rules could be "retired" after the launch event. On the other hand, new flight rules are written in order to incorporate lessons learned from flight experience gained via mission operations. See also the section entitled "Avoiding Human Errors."

SPACECRAFT POINTING CONTROL AND STABILITY PERFORMANCE^{1, 11, 23}

Spacecraft pointing control error is defined by the angle between the actual pointing direction and the desired pointing direction of a specific on-board body vector. The spacecraft pointing-control requirement is driven by the need to guarantee that the selected science target falls inside the field of view (FOV) of the science instrument. If the 2-mrad (radial 99%) pointing control requirement (see Table 1) for NAC is met, the captured image is guaranteed to fall inside the 6.1×6.1-mrad FOV of the NAC. As given in Reference

23, the flight performance of the Cassini spacecraft RWA-based pointing control and pointing knowledge are 0.639 and 0.629 mrad (radial 99%), respectively. Obviously, both the inertial pointing-control and pointing-knowledge requirements (listed in Table 1) are met with margins.

Spacecraft pointing stability is defined by the angle variation of the actual pointing direction of an on-board body vector over a time duration called “exposure time” (or “dwell time”). The spacecraft pointing stability requirement is driven by the need to ascertain that, over the exposure time of the imaging, incoming photons are “focused” on the intended set of camera charge-coupled device pixels. If the instrument moves during the exposure time, photons fall on that set of pixels as well as neighboring pixels, and a fuzzy image results. The specific pointing stability requirement selected corresponds to a tolerable degradation of image quality. Selected Cassini pointing stability requirements are also given in Table 1.²³

Sources of jitter inherent to the Cassini spacecraft include the reaction wheels’ static and dynamic imbalances, RWA bearing drag torque, and disturbance torque generated by sloshing motion of propellant in their tanks. Attitude determination sensor noise with a significant frequency content within the controller bandwidth looks like valid “commands” to the attitude controller. Accordingly, the controller generates control torque to cause the spacecraft attitude to follow these erroneous “commands.” This results in undesired spacecraft vibration.

In Reference 23, ten sets of “point and stare” science observation data were collected and analyzed to generate a range of Cassini’s pointing stability performance from 2003 through 2008 (prime mission was 2004–08). The dependency of the spacecraft pointing stability on exposure time, in log-log scale, is depicted in Figure 4. Superimposed on these plots is the pointing stability requirement mentioned in Table 1. These results overwhelmingly confirm that the Cassini pointing stability requirements are met with very significant margin. The quality of images returned by the high-resolution cameras provides ample evidence to this claim.




Figure 4. Summary of X-axis RWA-based Pointing Stability²³

RWA CONTROLLER DESIGN AND BEARING DRAG TORQUE ESTIMATOR^{11, 12, 24}

The reaction wheel assemblies are used primarily for attitude control when precise and stable pointing of a science instrument is required during the prime mission phase. The RWA controller consists of a bearing drag torque estimator in the FSW. With this drag torque compensation system, regardless of whether the bearing drag torque is at its nominal level or is elevated due to anomalous bearing performance, the appropriately compensated motor torque command will be sent. As a result, impacts of

the bearing drag on the S/C attitude control performance are minimized. This is one key reason why the Cassini pointing stability requirements are met with large margins as mentioned in the last section.

For Cassini, the drag torque estimator was designed to accurately track the bearing drag torque only in the steady state. When the physical drag torque changes due to, for example, a spin rate reversal (“zero-crossing”), the drag estimator can still track the physical drag torque but there will be transient tracking error. The faster the drag torque changes, the larger will be the tracking error.²⁴ As a result, the drag torque estimator will not be able to fully compensate for the physical drag, and the spacecraft attitude control and stability performance will suffer as a result. As an example, as depicted in Figure 5, there was a “zero crossing” of the RWA-4 rate at 2013-DOY-005T18:24:11 (as indicated by a bold arrow head). As a result of the incomplete RWA-4 drag torque compensation, noticeable perturbations of all three per-axis S/C’s attitude control errors were observed soon after that rate reversal event. For Cassini, this isn’t a problem because of the large performance margin (see the last section) and the transient nature of the performance degradation. However, for missions without the benefit of large performance margin, a more capable drag torque estimator that can better track transient drag “spikes” will be needed.

Since the year 2000, all Cassini RWA bearings experienced a class of anomalous drag torque that were generally “spiky” in nature.²⁴ The drag spikes usually occurred at a time when the RWA was maintained at a constant spin rate for a long period of time. In this spin condition, the expected level of RWA bearing drag torque is nearly constant. However, drag torque spikes were often observed superimposed on the “constant” drag torque. The initial sudden rise in drag torque was often time followed by either a rapid (several minutes) or gradual (several hours) exponential decay to the nominal drag level. The spikes have a wide range of magnitudes and they occurred in a wide range of RWA spin rate conditions.²⁴ A representative set of these RWA drag spikes is given in Figure 6. Again, the implemented Cassini RWA drag torque estimator could not fully compensate for these drag spikes, and the spacecraft attitude control and stability performance degraded slightly if the drag spikes are large. This is another reason to implement a more capable drag torque estimator that can better track transient drag “spikes”. More discussions on drag torque compensation control schemes are given in Ref. 66.




Figure 5. Transients of S/C Attitude Errors Caused by RWA-4 Rate Reversal on 2013-DOY-005.




Figure 6. Spiky drag torque observed at a RWA-1 spin rate of -600 rpm on 2006-DOY-038

REACTION WHEEL BIAS OPTIMIZATION TOOL (RBOT)¹²

In preparation for the extensive use of reaction wheels in the prime science mission (which started on July 1, 2004), the operations team began to develop and use a ground software tool named RBOT in early 2001. This software tool is used by the operation team to select an optimal set of RWA bias rates so that the total RWA dwell time inside the problematic sub-EHD (elasto-hydrodynamic) region can be minimized. As mentioned above, the use of reaction wheels for spacecraft attitude control is subject to three RWA speed constraints. The combination of these low and high-speed RWA spin rate constraints can present a significant challenge to the Cassini operation team on the use of RWAs to support a multitude of science slews over a long period of time. The ground software tool RBOT was developed by the Cassini operations team for the management of the Cassini RWA consumables.¹²

Given the predicted time histories of the spacecraft's attitude and attitude rate commands (to support science slews), from the start to the end of an RWA biasing time segment, RBOT will select a set of optimal biasing rates for the three prime RWA that minimizes a cost function. The cost function is designed to enforce the operational constraints described earlier (see Reference 12 on the definition of this function). However, for biasing segments with very complex science observations, it is possible that even the best set of RWA bias rates might still contain prolonged periods of spin rate dwelling inside the sub-EHD region. In this situation, one or more of the following remedial actions should be aggressively pursued to protect the wheels: 1) break the problematic biasing segment into two or more shorter segments; and 2) modify some problematic science observation sequence designs (e.g., slew the S/C using a set of slower rate and/or acceleration profile limits) to allow RBOT to find solutions without long dwelling inside the sub-EHD region. Disciplined and long-term use of RBOT has led to a significant reduction in the daily consumption rate of the RWA low-rpm dwell time. For example, the consumption rates of the three prime wheels' low-rpm times were 2.5 hours per day (each) in 2005–2006. By the year 2010–2013, the consumption rates were reduced to 0.35 hours per day (each). This drop in the low-rpm dwell time is due mainly to the disciplined use of RBOT.

Many spacecraft with attitude controlled by reaction wheels (or control moment gyroscopes) have encountered bearing-related flight anomalies. Over the past nineteen years, the Cassini mission control team has learned many useful lessons.²⁴ These lessons are listed below (in random order).

1. Trend RWA performance, beginning with wheel acceptance tests and throughout mission operations, to identify potential limitations on reaction wheel lifespan.
2. Implement a reaction wheel drag torque estimator in the flight software to provide ground visibility of any anomalous bearing drag conditions.
3. To assist in minimizing the time the wheels spend inside the sub-EHD region, develop and use a ground software tool (e.g., RBOT) to carefully manage RWA biasing events.
4. Aggressively and constantly look out for opportunities in science observation sequence designs that can reduce low-rpm RWA operations.
5. Besides the spin rates of the reaction wheels, probably the next most influential parameter that controls the performance of the bearing lubrication system is the bearing temperature. The bearing temperature must be carefully monitored.
6. Where RWA performance data indicates that the RWA lifespan may be constrained, take measures to mitigate the mission impact. For example, control the spacecraft's attitude using thrusters before the start of the prime mission and for observations with pointing stability requirements that could be easily achieved using thrusters.
7. Review the FP design to identify its vulnerability when RWA drag torque becomes elevated.
8. Design, test, and exercise contingency procedures that will be needed to recover the S/C from a safing state that is caused by a degraded/failed RWA.
9. If unambiguous bearing degradation trends are observed in multiple reaction wheels well before the end of the prime (or extended) mission, it is important to get ready with the implementation planning of a contingency hybrid controller (S/C attitude is controlled by a combination of thrusters and RWA).^{57, 72-73}

IMPACTS OF RWA INDUCED VIBRATIONS ON SCIENCE INSTRUMENT OPERATIONS

The Composite Infrared Spectrometer (CIRS) is one of the key Cassini remote sensing science instruments. In Figure 1, it is located just underneath the two star trackers. It captures infrared light and splits the light into its component wavelengths (or colors). It then measures the strength of the light at each of those wavelengths primarily to measure the temperature of objects, but also their composition. The moving scan mechanism subassembly of CIRS includes the optical and mechanical components in the optics assembly required for moving the interferometer mirrors to permit controlled sampling in the optical path difference. The scanning mirrors are normally maintained at a constant scan velocity. Strong external disturbances at the CIRS base plate can cause large scan velocity variations that disrupt the regular sampling process. This results in too many samples in the CIRS interferograms (raw data) that degrades the CIRS data. A key disturbance force comes from the static and dynamic imbalances of the three spinning reaction wheels. There are other microphonic sources on the spacecraft but only the attitude control RWA cannot be powered off at the time of CIRS operations.

Pre-launch, the impact of spacecraft vibrations on CIRS operations was a concern. Workshops on the S/C microphonic environment were conducted to address this concern. The evaluation of the effect of RWA induced vibration on the CIRS operation consists of three “components”: [1] the measured force output of the Cassini RWA over a range of spin rates, [2] measurement of the transfer function between the RWA force and the CIRS base vibration, as measured using the Cassini DTM (Developmental Test Model), and [3] the sensitivity of the CIRS scan velocity error due to the base vibration. In addition to this type of component-level evaluation, the RWA-to-CIRS end-to-end measurements had also been performed on the spacecraft.

Inflight, CIRS calibration data were found to be impacted by disturbances generated by RWA. This is the case in spite of the fact the static and dynamic imbalances of the reaction wheels are at least a factor of three smaller (that is, better) than their respective requirements. Correlation analyses between RWA's spin rate and the fraction of good CIRS calibration data indicated that low RWA rpm led to low data loss (that is, most calibration data are good). But when the RWA operated at high rpm (>1,300 rpm), calibration data loss become significant. The productivity of the CIRS science investigation was affected. Multiple options

have been employed to mitigate the data loss. In some science sequences, AACS team biased the RWA in order to operate them at rates lower than 1,300 rpm. This approach was only limited to cases when the resultant sub-EHD time of the wheel bearings isn't too long. An alternative remedy is to extend the duration of the CIRS calibration period in order to provide more data to compensate for those "contaminated" by RWA interference. These and other remedial options provided up to about 70–80% return of the CIRS calibration data. Missions with science instruments whose operations are sensitive to low-level vibrations should pay attention to microphonics generated by RWA imbalances. The evaluation of the RWA-to-instrument impact should be performed conservatively.

TUNING OF RCS ATTITUDE CONTROLLER PARAMETERS^{1, 10, 25}

After crossing 2.5 A.U. from the Sun, ground communication with the Cassini spacecraft was switched from LGA to HGA. This required tightening the RCS controller dead-band (DB) from [20, 20, 20] to [2, 2, 20] mrad about the S/C's X, Y, and Z-axis, respectively. This new deadband significantly increased the consumption rates of the two thruster consumables (on/off thruster cycles and hydrazine cost). Following the completion of the inner Solar cruise phase, the AACS team tuned the RCS controller parameters to improve these consumption rates ahead of the long outer Solar cruise phase.

During periods when fast changes in S/C attitude are not required, the RCS controller reverts from the conventional BOB control algorithm ("high rate" mode) to the Adaptive Pulse Width Adjuster mode (APWA, also known as "low rate" mode). The APWA logic is quite simple: The next thruster "on-time" is the last thruster "on-time" multiplied by a factor: $Factor = [\sqrt{L_2} + \sqrt{L_3}] / [\sqrt{L_1} + \sqrt{L_2}]$. Here, L_1 and L_2 are the "normalized" spacecraft angular excursions travelled by two previous one-sided dead-band excursions (see Figure 7). The normalization is made with respect to the two-sided dead-band. L_3 is the desired normalized excursion of the next cycle. With this logic, if L_1 is smaller than L_3 , then $Factor > 1$, and the next pulse width will be larger than the last pulse width, and vice versa. The size of L_3 was selected to achieve an angular excursion that is as large as possible but without "touching" the dead-band boundary. The expected benefits of changing L_3 was to reduce the number of two-sided dead-band motion and thereby reduce the average hydrazine consumption rate.

In the presence of a very small environmental torque that cannot reverse the polarity of the attitude rate (resulted from the last thruster firing) within the $\pm DB$ space, two-sided dead-band resulted. In this case, we select a multiplication factor of 0.6 to lower the thruster firing time for the next pulse. In addition, the commanded dead-band size is also increased slightly by a small step. That step is specified by the FSW parameter named "Walking Deadband Step" (WDS). A maximum of five "walking" steps outside the dead-band boundary are allowed. The expected effect of the changing WDS was to reduce the number of multiple pulses during two-sided dead-banding. This is important because for each thruster, the maximum allowable thruster on/off cycle for the entire mission is limited to less than 273,000 cycles.¹




Figure 7. Cassini Z-axis Attitude Control Error on 2001-DOY-204.

Two in-flight tunings were performed to determine the effectiveness of two new values of WDS and $\sqrt{L_3}$. The first test was performed from 2001-DOY140-T16:33:03 to 2001-DOY-144T12:22:54 with the WDS changed from 120 to 200 μrad . The second test was performed from 2001-DOY-203T16:03:30 to 2001-DOY-206T15:52:10 with $\sqrt{L_3}$ changed from $\sqrt{0.75}$ to $\sqrt{0.5}$. Test results were analyzed by both the AACS and Propulsion teams. Their analyses of data from the first test indicated a 19.3% reduction in the total number of thruster valve cycles and a corresponding 2.5% reduction in hydrazine consumption. For the second test, a 41% reduction in the hydrazine consumption was found.⁷ Based on these positive flight results, these new parameter values were made permanent in the AACS FSW.

The Cassini RCS controller design has a logic to autonomously switch between the “high rate” and “low-rate” RCS control modes. This high-to-low mode switch is performed if both the rate command and rate error signal are persistently low (<0.05 mrad/s) for a period of time T_{SWITCH} . Based on pre-launch simulation results, T_{SWITCH} was selected to be 10 min. Post-launch, T_{SWITCH} was lengthen to 40 min., resulting in reductions in the consumption rates of both propellant and thruster on/off cycle. This is another example on the need to revisit the values of controller parameter selected pre-launch which are often time done without a good knowledge of the environmental torque imparted on the spacecraft.

In conclusion, for mission with RCS thrusters having small minimum impulse bit, the RCS controller should be designed with an APWA-like logic to achieve single-sided deadband. Via inflight tuning, the logic could significantly reduce the consumption rates of two key thruster consumables (thruster on/off cycle and hydrazine cost) during the long RCS-controlled cruise phase of the mission.

ESTIMATION OF TITAN AND ENCELADUS ATMOSPHERIC DENSITIES USING ATTITUDE CONTROL DATA²⁶⁻²⁹

One of the major science objectives of the Cassini mission is an investigation of Titan’s atmosphere constituent abundance. To this end, the instrument Ion and Neutral Mass Spectrometer (INMS) plays a key role.^{1,30} Also, the Huygens Atmospheric Structure Instrument (HASI), mounted on the ESA Huygens probe, sampled and determined Titan’s atmosphere density (as a function of Titan-relative altitude) during the Probe’s 2.5-hour descent through Titan’s atmosphere on January 14, 2005. Post-launch, the AACS team devised a methodology to estimate Titan atmospheric density using RCS thruster data collected during low-altitude Titan flybys. These AACS-centric density estimates are used to calibrate and supplement density estimated using science instruments.³⁰

The Cassini spacecraft attitude is controlled by RCS thrusters during all low-altitude ($<1,300$ km) Titan flybys. Thrusters are fired to overcome environmental torque imparted on the spacecraft due to Titan atmosphere as well as to slew the spacecraft to meet the pointing needs of science instruments (such as INMS). Obviously, the denser Titan’s atmosphere is, the more thruster firings will be needed. In other words, thruster firing telemetry data could be used to estimate the three per-axis torques imparted on the spacecraft due to Titan’s atmosphere. Titan atmospheric density (as a function of altitude) could be estimated from the computed time-varying atmospheric torque if we know the spacecraft velocity relative to Titan, the projected area of the spacecraft on a plane perpendicular to the velocity vector, the offset distance between the S/C’s center of mass (c.m.) and center of pressure (c.p.), as well as the drag coefficient. Details of this methodology are given in Refs. 26–27. The uncertainty of Titan atmospheric density estimated using this methodology is related to the following error sources: the spacecraft projected area, the cp-to-cm offset distance, etc. The one-sigma estimated uncertainty of the computed atmospheric density is 6%.²⁶⁻²⁷

At the time this paper was being prepared, the methodology had been applied on attitude control telemetry data from 43 low-altitude Titan flybys executed in 2005–2013.²⁸ It had also been applied on data from 5 additional low-altitude Titan flybys executed in 2012–14.²⁹ These density estimates made using GNC data provided the Cassini project with a wealth of Titan atmospheric density data that was used to calibrate INMS-based density estimate.³⁰ This methodology is a significant contribution the AACS team has made to the attainment of a key Cassini science objective. A similar approach has also been used to estimate the Enceladus plume density.³¹⁻³³

ESTIMATION OF UNWANTED ΔV GENERATED BY THRUSTER FIRINGS

Figure 8 (from Reference 1) shows the locations of the four thruster pods that are mounted on a structure that is attached to the lower equipment module of the spacecraft base body. RCS thrusters are used to perform many AACS functions throughout the Cassini mission. During the cruise phase, they are used to maintain the attitude of the spacecraft, to slew the spacecraft, to detumble the spacecraft after it was separated from the launch vehicle, and to execute small ΔV burns. During the Tour phase, thrusters are used to bias the RWA momentum and to control the spacecraft attitude during low-altitude Titan flybys.¹ During these activities, Y-facing thrusters are always fired in pair. Hence, the forces generated by these Y-facing thrusters will cancel each other, and the ΔV imparted on the spacecraft will be negligible. On the other hand, the firings of Z-facing thrusters will generate unwanted ΔV on the spacecraft since they all point in the same direction. These unwanted ΔV will impact the spacecraft flight path and must be predicted by the AACS team for the Navigation team.




Figure 8. Cassini Thruster Pod Location and Thruster Orientations¹

The linear momenta imparted on the spacecraft due to the firings of the Z-facing thrusters could be estimated by multiplying the thruster magnitude with the on-time of the four Z-facing thrusters. However, thrusters do not respond “instantaneously” to their firing commands. There is an “on/off” delay time between the firing command and the time thrust is generated. Thereafter, thrust magnitude will increase with time exponentially (with a “rise” time constant) to its steady state value. Similarly, after the termination of the firing command, thrust magnitude will stay at its current value for another “on/off” delay time, followed by an exponential decay (with a “fall” time constant) to zero. The Cassini AACS flight software uses the calibrated magnitudes of the thruster’s force,^{34–35} on/off delay time, rise/fall time constants,⁸ and the estimated mass of the spacecraft to estimate unwanted ΔV . These parameters are routinely updated via ground command based on maneuver navigation reconstruction. The estimated thruster magnitudes and rise/fall time constants could also be trended to monitor the health of the thrusters.³⁵

In-flight, the AACS team uses a ground software tool named KPT (Kinematic Prediction Tool) to automate the prediction of ΔV (see also the section entitled “Ground Simulation and Tools”). The AACS team also uses FSDDS (Flight Software Development System) to predict the ΔV imparted on the S/C due to RWA biasing events, RWA drag torque characterization tests, as well as during low-altitude Titan flybys.^{36–38} The ΔV magnitudes predicted by the FSDDS are typically more accurate than those predicted by KPT because attitude control thruster firings are accounted for in FSDDS-based simulation (but not so in KPT). The requirement on the ΔV prediction accuracy, as imposed by the Navigation team on the AACS team, is ≤ 2.5 mm/s. Inflight, this ΔV prediction accuracy was usually met with significant margin – typical errors are ± 0.25 mm/s. But the requirement was violated five times due to ground software prediction limitations and one time due to a human error. The ISA (Incident-Surprise-Anomaly report) of these violations contain root causes that were studied and the ground software tools amended. Distribution statistics of AACS-centric ISA are given in the appendix.

ATTITUDE CONTROL SYSTEM FAULT PROTECTION DESIGN^{39–44}

There are two fundamental Cassini AACS FP requirements. Firstly, during all mission phases, the spacecraft must be able to “Fail Safe” by autonomously locating and isolating any single failure, recovering

to a thermally safe and command-able attitude, and then waiting for further instructions from the ground operators. Secondly, during a few time-critical events (including Launch,¹ SOI,^{16, 59} and Probe relay⁴⁶), the spacecraft cannot afford to simply isolate a failure and wait for ground instruction. In these “time-critical” events, the spacecraft must “Fail Operational” by autonomously recovering to a much larger set of its capabilities and then proceeding with a previously uplinked “critical” command sequence.

The architecture of the FP flight software algorithms that perform autonomous detection, isolation, and recovery from failures of AACS equipment and AACS-controlled propulsion elements is described in Reference 40. The front end of this architecture is a set of error monitors. Error monitors test the performance of AACS sensors, actuators, or functions (e.g., attitude estimation) against expectations. Deviations between the actual and expected performance are gauged against a pre-selected set of “thresholds.” An error monitor is “triggered” if its *threshold* is exceeded for a time duration that is longer than a *persistence limit*. Pre-launch, both the thresholds and persistence limits of all the error monitors were carefully selected by a team of FP engineers using ground-based test results. Flight experience indicated that most of these parameters were well selected but some thresholds and persistence limits must be changed in-flight because of anomalous sensor performance experienced post-launch or the presence of environmental torque.^{1, 41–45}

The Cassini AACS FP design contains more than four hundred error monitors. Not all error monitors are active during a particular spacecraft activity, but a large number of them are. The performance of these active error monitors must be monitored even if none of them are triggered. This is because one or more “un-triggered” error monitors might be a “hair line” away from being triggered. To monitor the performance of these error monitors, the “High Water Marks” (HWM) of the error monitors are computed by the FSW and sent down as telemetry. To efficiently gauge the performance of hundreds of error monitors, we first convert the HWM data into percentages of their respective thresholds. Computed values of these “percentages” are then displayed graphically (see Figure 9).^{1,41} One glance at this plot will provide the AACS FP engineer with a quick assessment of the FP performance. For example, from Figure 9, we see that the threshold of the error monitor named “IRU Parity Violation” (IPV) error monitor has been violated. Also, if the threshold of an error monitor was poorly selected pre-launch, one will see that the value of percentage quantity will always hover near 90–100%. A revision of the threshold might then be necessary.




Figure 9. In-flight Monitoring of the AACS Error Monitor Performance^{1,41}

Selections of error monitors' thresholds and persistence limits involved compromises. If the threshold is set too low, the monitor will be triggered even when there isn't any problem. This is a "false alarm" scenario. On the other hand, if the threshold is set too high, an abnormality might escape detection when a FP response is warranted. This is a "missed detection" scenario. In flight, the AACCS FP team uses the graphical tool described above to gauge whether the thresholds and persistence limits selected pre-launch represent good compromises. An example that is related to the error monitor IPV is given below.

Nominally, the prime IRU consists of three prime gyros and a "parity checking" gyro. With carefully selected weights, the weighted sum of these four gyros outputs (called parity) could be made very close to zero. If the parity sum is larger than a threshold for a time duration that is longer than a persistence limit, a fault is assumed to have affected at least one of these four gyros. In August 2000, the persistence limit was raised from 1 to 1.5 s in order to provide margin against false alarms due to occasional occurrences of gyro "spikes". Spikes are transient events triggered by the deposit of high-energy ions and/or protons on the gyro's input buffer. These transient events were not anticipated when the original persistence limit of IPV was selected. By raising the limit from 1 to 1.5 s, we have avoided "false alarms" of the IPV monitor. Similar "tuning" of fault protection parameters occurred several times inflight.⁴¹ Some of these changes are permanent, others are temporary. The effectiveness of the modified FP parameters must be monitored.

SPACECRAFT COMMANDING AND SAFE MODE⁴⁵

Cassini is designed to take care of itself, but to accomplish science and engineering activities the ground team operates the spacecraft using stored sequences of commands. Sequences are time-ordered sets of commands that are recognized by the onboard flight computers. Each command has a time-tag and that command is issued at that time. In practice, a sequence begins at a specific UTC (Coordinated Universal Time, onboard UTC is called SCET for Spacecraft Event Time) and each subsequent command is sequenced a given number of integer seconds after the previous command is issued. Commands are grouped into sequences that span many weeks. The ground team usually builds a background sequence about 5 months before it is uplinked to the spacecraft. Each background sequence spans about 10 weeks. If sequence S98 is executing onboard, sequence S99 is uplinked about one week prior to the end of S98 so that it is resident and ready to execute when S98 finishes. In parallel, the ground team is preparing sequence S100 and is planning S101.

Along with background sequences, mini-sequences and other real-time commands are designed for activities that cannot accommodate a 5-month lead time. These smaller sets of commands are also time-ordered and can run in parallel with an onboard background sequence. For example, the S98 background sequence may issue a command to turn the spacecraft to Earth-point at time T_0 . The turn takes 10 minutes, and when complete the spacecraft is tracking (pointing at) the Earth until another command is issued to begin a roll about the HGA axis at $T_0 + 30$ minutes. The roll finishes at $T_0 + 4$ hours. A ΔV maneuver mini-sequence is uplinked at $T_0 + 30$ minutes and reaches the spacecraft at $T_0 + 2$ hours. This mini-sequence will begin executing at $T_0 + 5$ hours by warming up the accelerometer. It commands main engine burn ignition at $T_0 + 6$ hours. Both the background sequence and the ΔV mini-sequence are running concurrently and care must be taken that they don't clash. An even smaller real-time command file could be uplinked at $T_0 + 1$ hour to specify downlink data rates after the maneuver is complete. All three sequences can run concurrently.

To meet the "Fail Safe" requirement, Cassini FP design will autonomously locate and isolate any single failure, and recover to a thermally safe and commandable attitude on thruster control. This process is commonly called a "safing" event. If safe mode is triggered, all stored sequences stop executing. Cassini autonomously turns to the "nominal" safe mode attitude.⁴⁵ An onboard "safe table" contains the vectors that define the safe mode attitude (along with an angular offset). This table can be changed via sequenced command. When Cassini was in the inner solar system, the nominal safe mode attitude was (see Table 2): NEG_Z to SUN (primary), POS_X to J2000Z (secondary). This insured both commandability through an onboard low-gain antenna and a thermally safe spacecraft.

Table 2. Examples of “Safe” Table Vectors that define the Safing Attitude

Safing Attitude	Primary Vector Pair	Secondary Vector Pair
Nominal Safe Attitude	NEG_Z to SUN	POS_X to J2000Z
HGA Safe Attitude	NEG_Z to EARTH	POS_X to J2000Z

In almost all cases, -Z (NEG_Z) to Sun does not permit downlink through the HGA (Earth-point is required). Therefore, a safe mode event required ground commanding even to achieve Earth-point. About 16 months before Saturn arrival, ground controllers uplinked new flight software that updated the on-board safe mode fault response. The goal was to try to minimize the complexity of reestablishing normal operations if a fault were to occur. The safe mode response will still begin by commanding the nominal Sun-pointed safe mode attitude, but after 60 minutes, fault protection algorithm will command a “High-Gain Antenna” safe mode attitude. By pointing the HGA towards Earth, engineering telemetry data will be receivable and this allows ground controllers to more rapidly assess and respond to the spacecraft anomaly that triggered the safe mode response. This HGA-safing algorithm in system fault protection has demonstrated its vital value several times in flight.

Each component of the Cassini AACS flight system defaults to a safe state upon any detected deviation, and recovery is aided by ample diagnostic information.³⁹ AFC bus communication watchdog timers are used as well as a heartbeat signal that AACS sends to the command and data computer every second. Were a heartbeat loss to occur, or a peripheral watchdog timer to expire, the AFC or peripheral would be commanded to reset. In these serious cases, the HGA-safing algorithm is inhibited – meaning that the spacecraft turns to Sun-point and stays there until ground controllers respond.

Both Sun and Earth-pointed safing attitudes need to be thermally safe and provide a clear star field for the SRU. J2000Z is the Earth’s spin axis at the January 1, 2000 epoch (the Z-axis of the J2000 inertial frame.) At Saturn, the clear star field preference is non-trivial: the rings, the planet, and bright moons require changes to the safe table depending on Cassini’s orbital inclination and proximity to these celestial bodies. The only cases where the HGA safing algorithm will not run are those that involve severe computer or severe attitude control faults. For those cases, the spacecraft will command the Sun-pointed attitude and ground controllers will not receive telemetry until they explicitly command an Earth-pointed attitude.

A complete list of safing events of the Cassini mission is given in Table 3. The “Cruise” phase of the Cassini mission was from launch (October 15, 1997) to December 31, 2003. In this 2,267-day period, the spacecraft attitude was mostly controlled by thrusters. Four safing events and an unexpected transition from RWA to RCS event were triggered during the cruise phase, wasting 909.52 gram of hydrazine. The “Tour” phase of the Cassini mission was from Jan 1, 2004 (start of the “approach science” phase) to September 15, 2017. In this 5,003-day period, the spacecraft attitude was mostly controlled by reaction wheels. Two safing events were triggered during the Tour phase, wasting 107.39 gram of hydrazine. Note that in Table 3, safing events 1, 2, 2a, and 4 were caused by anomalous AACS performance. Safing events 3, 5, and 6 were caused by anomalous CDS (Command Data System) performance.

Table 3. Cassini Safing Events, October 1997 – January 2016

Safing	Date	Hydrazine [g]	Root Cause(s)
1	1998-DOY-083	67.85	A swap of tracker triggered a “large Z/σ ratio” monitor.
2	1999-DOY-012	90.64	Slow roll tripped a “large Z/σ ratio” error monitor.
2a	2000-DOY-350	562	Low-rpm RWA operation led to high bearing drag. A triggered “Excessive RWA rate error” monitor caused a RWA-to-RCS transition <u>without calling safing</u> . Subsequent science mosaic slews on thrusters wasted significant hydrazine.
3	2001-DOY-130	106.70	Caused by a CDS configuration management error.
4	2003-DOY-132	82.33	Missing an IVT (Inertial Vector Table) target vector.
5	2007-DOY-254	68.18	SSPS (solid state power switch) of TWTA-B line was tripped off.
6	2010-DOY-307	39.21	File corruption causing swap from CDS-A to B.

Since Cassini launched in 1997, Cassini has put itself into safe mode a total of six times. Considering the complexity of demands the mission operation team has made on Cassini, the spacecraft has performed exceptionally well. This record supports the claim that the spacecraft is well designed and the Cassini Spacecraft Operations (SCO) team is well trained. Many members of the SCO team have been with the project since before launch, and their skills had been sharpened via Operational Readiness Tests (ORT) conducted by the SCO leadership. Their knowledge and training have enabled quick recoveries from the safe mode and to resume the mission.

ENGINE-BASED ΔV CONTROL AND OPERATIONAL ISSUES¹⁴⁻¹⁵

Delta V maneuver performance on Cassini has been outstanding. Since launch, 490 maneuvers have been planned. About 130 of those were cancelled as they were not needed. The NASA Deep Space Network (DSN) support has been excellent – not a single maneuver was missed and maneuver telemetry playback has been excellent. Maneuver design always includes a nominal (prime) design, and a backup design. The backup is performed if the prime maneuver does not execute, or very rarely, if there is a compelling navigation or operations reason why the backup maneuver is preferable. The backup is typically scheduled 24 hours after the prime maneuver. A typical maneuver DSN pass is 9 hours in duration with burn ignition planned 6 hours into the track. DSN maneuver tracks are thus scheduled in pairs – A 9-hour track for the prime, followed 24 hours later by a 9-hour backup track. A maneuver sequence is typically uplinked during the DSN track in which it will execute. This lets the ground team use the most up-to-date navigation data in designing the maneuver. In some cases, if orbit determination is stable and there is a significant ΔV penalty to fall to the backup, it is desirable to uplink the prime maneuver 1 or 2 days ahead of its execution time. This maximizes the chance of doing the prime maneuver and ensures that bad weather at a DSN station on the day of the maneuver will not affect its nominal execution.

Maneuvers are mini-sequences that overlay the Cassini background sequence. Ground rules are imposed on background sequence design so that maneuvers can be accommodated if needed but that can be seamlessly omitted too. For example, an RWA momentum bias may or may not be part of the maneuver sequence. But if a RWA bias needs to be guaranteed to occur, it is explicitly placed in the background sequence, independent of the maneuver. In those cases, it is placed during the maneuver window, but at a time after the maneuver is complete. The background RWA bias is placed 1.5 hours before the end of the 9-hour maneuver track. In some cases, the maneuver will command the RWAs to spin-down pre-burn, and spin back up post burn; the background bias then becomes redundant (this happens every ME maneuver on Cassini). In other cases, the RWA profile during a maneuver may be unfavorable (extended dwell time near zero-rpm, for example). Options include: (1) bracketing the maneuver with small RWA biases to improve the RWA profile; (2) placing a bias ahead of the maneuver, but having it match the background bias (which then becomes redundant).

The maneuver design process has to be rapid and robust. To achieve this, maneuver sequences are designed as “blocks” of time-ordered commands that are always issued. This provides great flexibility. Maneuver parameters are fed to the block including time parameters that permit variable-time activities. This ensures consistent commanding for maneuvers, and the expanded blocks can be carefully checked by the ground team prior to uplink. Most of the design occurs many days ahead of the maneuver, including early estimates by the navigation team of the orbit determination and burn ΔV vector. As the orbit determination is updated with more data, its effect on the maneuver (change in burn ΔV vector or turn times) can be evaluated and any problems mitigated. For example, an early OD (orbit determination) solution may produce a design that leads to an RWA over-speed issue. Options include slowing down the maneuver turns or bracketing the maneuver with RWA biases to avoid the problem.

A big effort was made to automate as much as possible the design and verification of maneuver sequences. The Maneuver Automation Software (MAS) ground tool is the result of this effort.⁵⁰ In less than 30 minutes, a complete maneuver sequence can be designed and verified using MAS. This tool ties many tools together. It runs on the Cassini operational computer network, but communicates directly with the navigation network too. Orbit determination and navigation maneuver design tools are run first. Then attitude control tools construct the vector and burn parameters needed to point the spacecraft and command the burn. Then sequence generation tools, known as SEQGEN and SEQTRAN, are automatically run to expand the maneuver block, check it, and generate uplinkable products. Then an extensive series of checks

are performed on the maneuver products, including kinematic, thermal, and RWA momentum predicts to ensure correctness and compliance with all flight rules. Commands to suspend star identification (if needed) due to Saturn, rings, or other bright bodies are merged to create the final products.

One important lesson from Cassini maneuver design is that over the course of the mission a large number of changes occurred that affected the basic design of the maneuver sequence. Had maneuver design and implementation been placed in onboard flight software, this would have caused major problems. Ground tools are much easier to update and validate than flight software. Changes such as swapping to the backup thruster branch, allowing longer settling time after maneuver completion for flexible magnetometer and propellant slosh modes to dampen out,⁶⁴ adding a 0.9° pointing offset to account for an unobservable mounting error in the rocket engine assembly – all these would have caused substantial flight software changes if the process were mostly onboard. The ground maneuver block is configuration controlled, but can be updated much more easily than flight software.

Burn execution error is a measure of how closely the spacecraft achieves the desired maneuver ΔV . The ability to accurately achieve the desired ΔV is important to Cassini mission operations and planning. For example, if ME burn ΔV execution was only accurate to 5%, the mission would have used up all of its bi-propellant years ago. If burn execution was perfect, there would be more propellant available and the mission might have lasted beyond 2017. Since ΔV is a vector quantity, burn execution error has a magnitude and a direction (see Table 1). The magnitude error is referred to as an under-burn if the achieved ΔV is less than the desired ΔV . An over-burn means the achieved ΔV is greater than the desired ΔV .

For ME burns, a key sensor on Cassini is the single-axis (spacecraft Z-axis) ACC. Accumulated ΔV from this sensor has been used for each one of the 183 main engines burns since launch. Each burn has cutoff at the targeted ΔV magnitude. A back-up timer, normally set at 5% above the nominal burn duration, has never been triggered in flight. The Cassini ACC bias has been stable (and is calibrated prior to every ME burn) throughout the 19⁺ years of the Cassini mission. The ACC bias showed more variation in the early days of the mission, perhaps due to solar heating effects when Cassini was traveling through the inner solar system.

The ACC scale factor is key in main engine burn magnitude execution error. Calibrating the scale factor is done on the ground using Doppler-derived execution errors over many in-flight burns, and any adjustment is carefully considered and tested before an on-board update (which requires a FSW patch). The Cassini ACC scale factor has been updated four times in flight. The first update after TCM-1 was a 1% change to match the ground-calibrated value. Between 2000 and 2009, three updates, each 0.06% or less, were made to the ACC scale factor based on reconstructions of the most recent navigation Doppler flight data to that point.⁶¹ In 2009, an adjustment of 3 mm/s to another FSW parameter (ME tail-off impulse) was made in the ME cutoff logic due to a consistent ME under-burn. Since 2009, the scale factor and its onboard estimate has been very steady and now the biggest source of uncertainty in ME burn magnitude is the amount of thruster firings that occur following ME burn cutoff. Following maneuver cutoff, attitude control switches back from TVC to RCS control. In some cases, cutoff occurs while the attitude control error is above the RCS “deadband” threshold, leading to more thruster firings as the RCS controller damps out the error after cutoff. Thruster firings during this period can vary from 2 to 7 mm/s of additional ΔV post-cutoff, and the random magnitude cannot be accurately predicted a priori.

Early in the mission, an error in the FSW ME ΔV magnitude logic was detected thanks to high-fidelity ground simulations of ME burns. In preparing for the second ME burn of the mission, simulations showed a 1% underburn in the 450 m/s Deep Space Maneuver (DSM). Originally this was thought to be an error in flight software estimation of the ΔV magnitude, but FSDS simulations also showed the 1% error in the “physics” or “true environment” accumulated ΔV . The FSW error was found in the ACC-sensing-axis to thrust-vector “de-projection” (the ACC sensing axis is offset from the thrust vector by about 7°) and fixed via a FSW patch prior to the DSM.¹³

Spacecraft pointing during the burn is controlled by gimbaling of the main engine. Two linear engine gimbal actuators provide two-axis thrust vector control with the third axis (roll about the thrust vector) controlled by RCS thrusters. The main engine pre-aim thrust vector is a fixed but updateable vector in the spacecraft body frame. This body vector is used to pre-position the engine gimbals towards the approximate center-of-mass of the spacecraft. The commanded spacecraft burn attitude is defined by aligning the pre-aim vector with an inertially-fixed desired ΔV vector defined by the navigation team.

Main engine pointing accuracy can be enhanced by carefully selecting the ME pre-aim vector. At launch, the best estimate for the pre-aim was simply the vector connecting the engine gimbal with the center-of-mass of the spacecraft. But because knowledge of the center-of-mass in the S/C frame is uncertain, in practice the best method of selecting the pre-aim vector for the next ME burn is to use thrust vector telemetry from the most recent ME burn. One complication is that the ME burn must be long enough to provide at least one complete cycle in the 0.03–0.04 Hz sustained limit cycle in the thrust vector believed due to the interaction between the TVC control action and the reaction flexing of the soft mounts against the propellant line.¹

Using thrust vector telemetry for the pre-aim vector works well from burn to burn but in some cases a correction should be made if a sudden change in the center-of-mass location occurs before the next ME burn. This has happened most notably in flight when the 11-meter magnetometer boom was deployed prior to Earth swingby, and also when the 320 kg Huygens Titan probe was ejected from Cassini on December 25, 2004. Reference 14 explains how this correction is made. This correction is also applied after many small ME burns occur, none long enough to provide adequate telemetry of the thrust vector over one 0.04 Hz oscillation. Most pre-aim updates adjust the vector by approximately 1 to 2 mrad (the shift was 4° after Huygens Probe release).

Main engine execution error (both magnitude and pointing) has been trended throughout the mission with Reference 15 providing plots based on flight reconstructions and are still accurate. Post-launch, the Cassini ME ΔV execution accuracy (both magnitude and pointing) has improved steadily with flight experience.⁶¹ For example, the Gates proportional pointing accuracy of ME burns improved from 10 mrad (1 σ , see Table 1) pre-launch to 1 mrad by the time of OTM-326 (after 145 ME ΔV burns). The corresponding Gates proportional magnitude accuracy of ME burns improved from 0.35% (1 σ , see Table 1) pre-launch to 0.02%.⁶¹

THRUSTER-BASED ΔV CONTROL AND OPERATIONAL ISSUES

RCS ΔV maneuvers are performed when the desired ΔV is less than what the main engine system can safely perform. At launch, this ΔV “cutoff” point was chosen to be about 0.5 m/s, but during Saturn operations ME burns were performed with ΔV targets as low as 0.27 m/s in order to conserve hydrazine. The shortest allowable ME burn duration, consistent with hardware requirements, is 1.0 second. The shortest ME burn actually performed was OTM-93 with a 0.27 m/s ΔV and a burn duration of 1.5 seconds. Now that the mission is reaching its conclusion with healthy B-branch thrusters and more than adequate hydrazine reserves, the cutoff is being moved back to near 0.5 m/s (to reserve the very limited ME bi-propellant for large burns).

RCS maneuvers are performed using the four 1-N -Z-facing thrusters. During the latter stages of the mission at Saturn, each thruster produces a force of about 0.6 N. The commanded spacecraft burn attitude for RCS ΔV maneuvers is defined by aligning the S/C -Z axis with an inertially-fixed desired ΔV vector defined by the navigation team. Since the accelerometer was designed for ME burns, the quantization of ACC ΔV measurements is not fine enough (1 ACC count equals about 4.5 Ns of impulse) to allow accurate estimates of ΔV during RCS burns. Instead, a “virtual accelerometer” is used which estimates thrust acceleration for each firing Z-thruster based on its on/off status, and on-board estimates of spacecraft mass and the nominal thrust for each thruster (parameters in the maneuver mini-sequence). From these quantities, the FSW computes an estimated acceleration for each firing thruster. In an RCS burn, all four Z-axis-facing thrusters are nominally firing together. Burn cutoff is commanded when the accumulated ΔV magnitude reaches the target. Without an explicit sensor, RCS burns are essentially “timed” burns where total thruster on-time is accumulated. So, ΔV magnitude execution error is directly related to the FSW estimates of the thruster forces. Ground tools estimate these forces using tank pressures and related propulsion system data,⁷ but the best way to keep the onboard settings accurate is to adjust them based on recent navigation (e.g. Doppler) flight data.

In 2008, the prime thruster A-branch experienced significant thrust force degradation which caused 2 of the 4 -Z-facing thrusters to lose 25-40% of their nominal thrust. This led to a 5% ΔV under-burn during a significant RCS ΔV maneuver. This is one limitation of timed burns. Using sensed acceleration in the onboard ΔV algorithm is the best way to avoid this, assuming a less-granular ACC onboard. Early in 2009, the ground team switched to the backup B-branch thrusters.⁵⁸ As of December 2016, no significant

degradation has been observed on the B-branch. The cause of the A-branch degradation is not definitely known.⁵¹ Theories include catalyst bed contamination or unexpected degradation due to many years of minimum impulse bit attitude control followed by intensive thruster usage after arrival at Saturn.

RCS ΔV pointing execution errors are mainly a function of how close Cassini's actual attitude follows the commanded burn attitude during the course of the burn. Cassini turns to the burn attitude using reaction wheels, so the pointing error at burn ignition is essentially zero. Just before ignition, a transition to RCS control occurs and the burn executes in RCS control while the reaction wheel momentum is held constant throughout the burn. Attitude control deadbands of 0.5° for the X and Y axes and 1° for the Z-axis limit the maximum attitude control error during and after the RCS burn. Since the center of mass is offset from the spacecraft centerline, 4 -Z thrusters firing at once do induce a torque that causes the attitude error to migrate to one side of the attitude control deadband. When an X or Y deadband is reached, 2 of the 4 -Z thrusters pulse off to provide torque to stop the growth of the attitude control error in that axis (if both deadbands are reached at the same time, 1 of the 4 -Z-thrusters pulses off). If the off-pulsing was just long enough to stop the growth of the attitude error, the attitude error would tend to chatter at the attitude control deadband ("ride the deadband") during the burn. This is undesirable both for the thruster hardware itself (too many pulses) and would lead to the thrust vector staying offset from the desired burn attitude by 0.5° in the X and Y axes. This offset could be accounted for by biasing the initial attitude, but many pulses not only degrade the thruster hardware but introduce ΔV magnitude errors due to uncertainty in the rise time and tailoff impulse of each thruster pulse.

To avoid these problems, a pulse adjuster in the FSW is used to compute off-pulse duration. Different implementations are possible. One would be "phase plane" logic (commanding off-pulse durations based on a combination of attitude and attitude rate). On Cassini, an attitude error integrator or "summer" is used. The attitude control error during the burn is integrated and multiplied by a gain. This signal is used to lengthen the off-pulse duration in each spacecraft body axis. The larger the accumulated attitude error, the longer the off-pulse. The values of these integrators are carried over from one RCS burn to the next. This helps subsequent RCS burn avoid "riding the deadband" during the burn. This results in better pointing, longer off-pulses, and fewer total off-pulses. A disadvantage of the summer approach is that the summers start at zero and RCS burns early in the mission did tend to chatter near the deadband. It took a number of burns to charge up the summers to the point that the pulse-adjuster was able reduce the chattering and the pointing execution error. Also, flight computer resets can cause the summers to go back to zero, requiring a patch to return them to their earlier values.

The thruster off-pulsing is accounted for in the ΔV accumulator algorithm. Typically, off-pulsing accounts for about 10% of the burn duration (i.e. the thrusters are on about 90% of the time during the burn). Individual thrusters typically average between 80% and 100% "duty-cycle" on-time during RCS burns. The duty-cycles tend to be smaller the further away the center of mass shifts from the Z-axis of the spacecraft. The backup maximum burn timer must be selected including the effects of the off-pulsing. To date, all RCS ΔV maneuvers have cutoff at the accumulated (virtual) ΔV magnitude, rather than at the maximum burn timer.

Post-launch, the Cassini RCS ΔV execution accuracy (both magnitude and pointing) has improved steadily with flight experience.⁶¹ For example, the Gates proportional pointing accuracy of RCS burns improved from 12 mrad (1 σ , see Table 1) pre-launch to 4.5 mrad by the time of OTM-328 (after 96 RCS ΔV burns). The corresponding Gates proportional magnitude accuracy of ME burns improved from 2% (1 σ , see Table 1) pre-launch to 0.4%.⁶¹

OPERATIONAL LESSONS LEARNED: AVOIDING HUMAN ERROR⁵²

Avoiding human error in spacecraft operations is a vital aspect of mission success, but is often overlooked when a mission's detailed requirements and capabilities are established. Many space missions have been degraded, or ended, due to human error during operations. Cassini is beginning its 20th and final year in space and this year Cassini will fly inside the rings so the operations team is especially cognizant of the need to minimize the effects of human error. In any long duration mission, ground engineers or controllers will sometimes make mistakes, perhaps just through complacency. The goal is to establish a mission operations system – meaning an operations team and their tools and processes – that identifies and corrects mistakes before any incorrect commanding reaches the spacecraft.

The Cassini spacecraft was designed to be fault-tolerant and robust and great care was taken in building the spacecraft so that the computers, propellant tanks and lines, sensors, and actuators would last for at least a decade, and hopefully much longer. The design of the AACS flight software also aids operators: for example, it will reject any command to power off a required resource (for example, the prime IRU). The design also provides good visibility of spacecraft states. This is key in regular operations and is especially important if an anomaly causes autonomous fault protection action. On Cassini, a large dedicated buffer in the AFC, called the fault protection event log, contains the entire history, with time tags, of any onboard fault response. This kind of visibility greatly aids in understanding the root cause of a spacecraft anomaly.

Some mistakes occur because the ground controller did not follow a procedure correctly. Others occur because there is no documented procedure at all. Sometimes, an experienced engineer will go on vacation and not have an adequately trained backup. Other times, a new engineer will come on board and make a mistake that no one detects. Adequate training, having an experienced backup to check all flight products, regularly scheduled peer reviews of each command sequence product, and not rushing through a verification process to meet a deadline, all these and more have to be part of standard operating procedure. Attention to detail, flying what you test and testing what you are about to fly, these are not just truisms but are the only real way to successfully operate a complex spacecraft far from Earth.

An important way to track down and fix a human error is to test a spacecraft command sequence on the ground before it is uplinked to the spacecraft. A ground simulator, whether it is a laboratory with actual flight computers and avionics,^{53, 59-60} or a pure software simulation,³⁶⁻³⁸ will normally detect an error in a command sequence. The authors are aware of two missions in which tests were performed and violations in the form of plots were generated, but ground operators missed the violation because the simulation outputs did not clearly highlight the violations. The sequence was uplinked and the spacecraft went into safe mode as a result. Bottom line: make sure violations are clearly flagged, preferably in a violations report. And these violations reports must be respected: every entry in it must be either waived (with a clear reason) or fixed.

There are innumerable ways human error can creep into ground control operations, so the Cassini team keeps an updated database of anomalies for the life of the mission – from early in flight software and hardware development pre-launch, all the way through the end of the mission. This database is a JPL-wide resource and documents not just flight anomalies, but operations and ground software anomalies too. If a ground simulator detects a violation, but an engineer misreads it, the problem needs to be documented in the ground software anomaly database and corrected in a timely fashion. The anomaly should not be closed unless some direct action is taken to rectify it.

Basic to any spacecraft operational team is the need to maintain a team of experienced engineers to operate Cassini. Operations engineers are often categorized differently than design engineers. Designers are crucial to putting together a quality spacecraft that is durable. Operations engineers are usually not experts in the design and may have limited understanding of flight hardware or software details. Design engineers tend to work with cutting-edge technology where new designs or approaches are encouraged, while operations people tend to “stick with what works”. Designers are often labeled more “creative”, while operators are sometimes considered more “turn the crank” users with limited technical understanding.

However, real world experience suggests that actual designs can have many “traps” unless they are carefully informed by operational knowledge. A design that is easily broken or mishandled is a poor design. In spacecraft operations, it is the operators who come to know the actual performance of flight hardware and software. Operators on some space missions have to build a database of “idiosyncrasies” and troubleshoot around unwieldy designs. It can be hard to predict the behavior of some designs. A few spacecraft have onboard intelligent agent flight software that might override a commanded slew path, for example. But if real world experience demonstrates that the behavior of the agent cannot be routinely predicted, it may be of little value and can actually detract from stable operation.

The AACS operations team provides formal and informal tutorials to new AACS engineers to help them learn about the flight hardware and software, how the spacecraft is commanded and what telemetry is produced on a daily basis. A new analyst becomes acquainted with key ground software tools and processes to build and validate command sequences. She supports many meetings relating to science and engineering activities for a specific background sequence. He supports meetings on OTM preparation, approval and post-burn wrap-up. She monitors the actual OTM burn to develop a familiarity with real-time telemetry

monitoring. He learns about ΔV maneuver execution error and how flight software parameter updates help minimize it. She learns under what conditions a fault protection response in the onboard flight software can become activated and autonomously issue commands onboard to recover after a hardware fault. He learns about how RWA health is optimized by judicious selection of RWA momentum commands at key intervals during a sequence.

A new AACS analyst learns about a key method of avoiding operational mistakes on Cassini – flight rules. Flight rules are operational limitations imposed by the spacecraft system design, hardware and software, violation of which would possibly result in spacecraft damage, loss of consumables, loss of mission objectives, loss and/or degradation of science, and less than optimal performance.²¹ Flight rules are an important way to capture the knowledge of both designers and lessons learned from previous operations on Cassini. Flight operations engineers need to become steeped in the mission flight rules. They need to know not only what the flight rules are and why they exist, but also how they are checked.

Human error is more likely when flight operations engineers do not have adequate tools. This is discussed more in the next section. Besides having fully capable tools, human error is more likely to creep in when an engineer:

1. Is not given enough time to complete the given task.
2. Does not adequately understand the task and its impact.
3. Does not take enough care in setting up or preparing for their task.
4. Does not follow a clear procedure.
5. Does not review his/her verification results carefully enough, or with others.

An example of knowledge deficit occurred when an AACS engineer designed an IRU calibration about one year after Cassini reached Saturn. To precisely estimate gyro scale factor errors during slews, good attitude estimation from star identification is essential. The AACS analyst mistakenly designed the slews for the calibration in a way that caused Saturn and Rings bright bodies to enter the SRU field of view. This caused star identification to be suspended and led to an inadequate gyro calibration.

A new flight rule was added to avoid the same mistake happening again. At the same time, the rule was explicitly added to ground software so an automatic check would detect the error in the future. Peer reviews were added as a standard review step for all future sequences. At the peer review, the entire AACS team (including cognizant FSW and fault protection leads) reviews all the verification results that an AACS lead analyst has done. The lead systems engineer for the sequence also participates. The lead AACS analyst reviews all his inputs and analysis for the sequence with the rest of the AACS team, including the AACS team lead. The informal nature of the meeting – about 6 to 8 people participate – is conducive to asking questions and allowing each engineer to learn in detail about issues that inevitably are relevant to their own lead responsibilities too. Peer reviews definitely help in finding errors.

Written procedures to document all the steps that an analyst should follow throughout the sequence development and verification process have been a great aid. On Cassini, procedures are living documents – currently kept on a Wiki – that are amended and clarified so that new analysts have a roadmap to help them follow the numerous steps needed to fully verify the command sequence.

Since absences or sickness can happen anytime, fully capable “backup” engineers shadow the lead AACS engineer for all background and OTM sequence design and verification. These backups are a crucial additional set of “eyes” because everyone at some point overlooks a parameter value or might miss a procedural step. Backup engineers independently perform the sequence verification procedure as a check on the lead analyst’s work. The backup engineer is especially important when the lead AACS analyst for a sequence is relatively inexperienced.

The Cassini team also uses detailed checklists for sequence verification that included sanity checks along with formal flight rules. A different checklist is used for OTMs or other real-time activities. Certain unique steps are required during final verification and a checklist ensures that each analyst remembers these key steps. Both the lead and backup engineers should participate in checklist completion. These checklists evolved from previous experience where an analyst may have made a mistake, so checklists especially alert an analyst to areas of potential human error that require extra care.

Important activities benefit from ORT. These involve the full spacecraft engineering team, often the navigation team, and science and uplink teams. One or more faults are introduced into the Cassini integrated test laboratory (ITL) -- which represents the spacecraft in this case -- which is configured for the

activity, and the team must figure out how to deal with the anomaly, often with certain additional “surprises” thrown in. Perhaps a hardware anomaly is introduced and downlink is lost. Or one of more key engineers is removed from the operational process so that other engineers need to demonstrate their own abilities to safely recover. Each ORT has a “lessons learned” review so that the flight team’s performance is critiqued and processes improved. ORTs are great for training new team members too.

Command approval meetings need to be scheduled so that all teams have adequate time to evaluate the actual flight products. Avoid any changes in the command files after verification and before uplink. Standardize steps like command file-creation-time review and simple file naming conventions to ensure there is no last-minute confusion about the specific command files to be uplinked.

GROUND SIMULATION AND TOOLS – KEYS TO FINDING ERRORS

Cassini utilizes a hardware integrated test laboratory for both flight hardware and flight software development and test^{36–38, 53, 59–60}. Flight spares of many flight hardware components provide outstanding fidelity and make the Cassini ITL an essential operational tool after launch. When combined with robust ground support equipment, the Cassini ITL is the perfect platform to test to-be-flown command sequences. The input products to ITL should be the actual products that will be uplinked, where possible. The ground telemetry system was fully integrated into ITL so that ground operators get a preview of their visibility of the actual activities on the spacecraft.

Cassini tests all critical and first-time events in ITL. ITL telemetry is easily accessible both during and after each test. ITL allows access to environment “truth” simulation parameters as well as FSW memory locations and telemetry. Besides telemetry, standard ITL output products include fault protection event logs, and records of detailed inputs used in simulation initialization.

Cassini always schedules ITL procedure walk-throughs so that both the ITL test team and the flight operations team have a common reference to follow for test initialization. Test initialization is especially important because an effective test must match the states of the actual spacecraft. This is especially true for flight software fidelity and as many subsystem parameters as possible. Test review must be a team responsibility and flight team reports on the outcome of each test need to supplement the ITL team output products. *A test not carefully evaluated provides little value-added.*

The ITL is also a great resource for in-flight anomaly evaluation too. To supplement the ITL, the Flight Software Development System (FSDS) is an “all software” simulator that was used to build and check out the flight software pre-launch.^{36–38} It is used extensively in operations because it contains all the flight software (including fault protection), runs faster than real-time, and runs on an analyst’s workstation. FSDS is used for detailed anomaly or engineering investigations, OTM testing, hardware-related fault insertion, and flight software development and regression testing. FSDS also contains updated models of Saturn and Titan atmospheres and Enceladus plume models. *A key aspect of operational success on Cassini is the careful design of test cases using FSDS or ITL.* Testing that emphasizes actual operational scenarios is especially valuable in tracking down human errors that could occur in flight.

Ground software (GSW) and procedures to verify command sequences, check flight rules, design science observations, and design ΔV maneuvers and other engineering activities are imperative for insuring errors do not get into uplinkable products and ultimately to the spacecraft. It cannot be emphasized enough that directing resources to GSW throughout the lifetime of a mission is key to safe operations. To date, over 160,000 spacecraft turns to point Cassini to accomplish science observations have occurred since reaching Saturn. Virtually all of these slews have been designed by the science teams themselves. Integrating all these activities together and insuring they meet all applicable flight rules and constraints requires robust GSW that both science observation designers and the engineering operations team utilize. The integration effort is huge and no hardware integration lab could possibly test more than a small percentage of these activities.

Spacecraft commanding requires a project-wide command database and a ground software tool used by all teams to generate commands for the spacecraft. At JPL, a tool called SEQGEN is used and the Cassini adaptation incorporates extensive flight rule checking as well as syntax and range checking on all command parameters. Outputs of this tool include a violations file and a “restricted” command summary. Certain commands are considered especially sensitive in that an incorrect argument could have dire consequences

to spacecraft health and safety. Extra visual inspection must be done for any restricted command including physically signing a form allowing flight usage.

The bulk of ground software development on Cassini occurred pre-launch and during the 6.7-year cruise from Earth to Saturn. Two key AACS ground software tools are the IVP tool and the Kinematic Predictor Tool (KPT). IVP generates all the fixed and time-varying vector commands that populate the onboard inertial and body vector tables during sequence execution.¹⁹ The KPT processes all the pointing-relating commands in a sequence, models all the pointing (slewing to and tracking targets) throughout an entire sequence, and models all flight rules that relate to spacecraft attitude (including the position of the sun and other celestial bodies). KPT emulates key AACS flight software but runs about 400 times real-time, allowing rapid verification of the 10-week background sequence.²¹

Significant resources were devoted to validation and verification of each ground software tool. This verification effort was especially intense pre-launch and during the 6.7-year cruise to Saturn. Even during the orbital mission at Saturn, many tools required extensive updates or fixes. Each ground software tool is configuration-controlled by the project and upgrades and fixes require project manager approval. As an example, KPT and IVP ground software went through many iterations to ensure vector and turn modeling was completely consistent with actual spacecraft behavior, even for unexpected events like having one turn interrupt another. KPT implements the SRU bright body flight rules and creates the sequence commands to suspend star identification when needed.⁵⁴ It also creates the commands to clear the AACS fault protection high water marks at the end of every downlink pass. IVP development continued during the Saturn Tour to enhance several science target observation categories, such as inertial vector definition commands for Radio Science observations requiring especially accurate time-varying pointing. KPT development continued during the Tour to enhance science instrument thermal modeling and improved ΔV predictions for RWA momentum changes. As recently as 2009, a new feature called Y-Biasing was added to KPT to allow slewing to an attitude tailored for that RWA momentum bias to reduce hydrazine usage and use the Y-firing thrusters more, saving the Z-thrusters for ΔV maneuvers and Titan flybys.⁵⁵

An important element of ground software is clear visibility of all flagged violations. KPT produces a violation summary report in a tabular format. This report documents all violations, when they occur, and list the science request in effect at the time of the violation. The KPT output log file is also important. This file should document all input and output files as well as include clear details about what caused each violation. One AACS analyst ignored a clear violation because the log file did not show him the input arguments of the command that caused the violation. The KPT log file needs to give the analyst as much visibility as possible into what led to each violation.

Engineers must have confidence in the ground software tools they use. Configuration control is vital, especially making sure that outdated versions of ground software cannot be accidentally invoked by analysts. Just as important is to ensure that ground software does not incorrectly flag a violation. Real violations may be ignored if an analyst notes that the tool is flagging violations incorrectly.

One aspect of turn modeling that might introduce an error is a turn that is near a singularity. On Cassini, the flight software will choose the shortest path to achieve a desired 3-axis spacecraft attitude. When this turn angle approaches 180°, a tiny difference in attitude can result in a turn in the direction opposite from what was expected. This occurred during a gyro calibration on Cassini in 2003 and led to trouble because the actual path chosen in flight required autonomous evasive action because of a boresight-to-Sun constraint violation.^{1,20} A flight rule was written to ensure that turns near 180° were flagged as violations so that workarounds can be identified during sequence design.

AVOID CHANGING FLIGHT SOFTWARE BY PATCHING WHEREVER POSSIBLE

One way human error in operations can adversely affects a spacecraft is by changing flight software without sufficient rigor. Mars Global Surveyor encountered this in 2006.⁵⁶ Flight software configuration management on Cassini has been strictly followed throughout the mission. A formal process involving project manager approval is required to even begin the effort to make a change to the flight software source code. Then extensive ITL and FSDS regression testing must demonstrate that the change has no adverse impact on any operational scenario. Testing must clearly show the “before” and “after” effect of any change. Then procedures and command sequences must be generated and successfully executed in ITL to demonstrate that the process of uplinking and loading the FSW can be seamlessly integrated with the

onboard background sequence. Only after all these checks are passed does the project manager permit the flight software change to be made onboard.

The AACS FSW that was used at Launch was A6.3.5. Consistent with the plan made at the time of launch, the AACS FSW was updated with new versions several times during the long cruise phase and during the early Tour phase. A summary of all full-load FSW updates is given in Table 4.²²

Table 4. Cassini AACS Full-load FSW Updates and Patches²²

FSW	Uploaded on	Key FSW capabilities added
A6.3.5	Oct 15, 97	Launch load
A7.7.6	Mar 7, 00	Detect thruster leak during long cruise; ⁴² enhanced RWA control performance; ¹¹ enhanced SID algorithm to handle extended bodies in tracker's FOV; "deluxe" attitude initialization capability (for SOI); IVP with rotating coordinates (Probe tracking); gyro-less attitude estimation capability.
A8.6.5	Feb 16, 03	This is the load used by TCM-19b as an inflight confirmation of the "energy-based" burn termination algorithm (for SOI). ¹⁶
A8.6.7	Apr 27, 04	This is the load for SOI burn. All SOI-related RCS slews were checked with this load. Confirmations of FSW capability to autonomously disable six SSA error monitors using the eclipse command, and to propagate attitude with long-duration SID suspend.
A8.7.1	Oct 2, 04	This is the load for Probe release and probe relay. ⁴⁶ Parameter updates for Probe release and relay. Reference trajectory updates for Probe relay.
A8.8.0	Oct 22, 10	FSW changes include RWA safing design, fixes to telemetry channels, updated thruster magnitudes, tracker misalignment matrix, AACS recovery data, and active RAM patches.
A8.9.0	Dec 12, 12	FP design to handle B-thruster leak in a mixed-branch RCS control, updates to default Safing attitude, default thruster magnitude, mass properties, CMT acceleration limit, and active RAM patches. This is the last FSW upload.

In Table 4, both A8.7.1 and A8.8.0 were full FSW loads. In between them, there were five patch loads: A8.7.2 (May 05), A8.7.4 (April 06), A8.7.5 (January 07), A8.7.6 (January 08), and A8.7.7 (June 09).⁴³ The full load A8.8.0 was based on patches introduced by A8.7.1 to A8.7.7. Human error can occur if a small "patch" to the flight software is introduced. The normal way ground controllers command the AACS flight computer is via goal-oriented sequence commands. Sequence commands have built-in checks to ensure commands are not corrupted. They also have mode and range checks. If an operator makes an error in a command argument, the flight computer will reject the command. Memory write commands ("patches") bypass most of these checks. That is why "patching" via memory write is strongly discouraged on Cassini.

In the rare case that a patch is required, the AACS team must follow the same rigorous process of testing and project approval as for a full flight software source upload. Cassini experience has shown that any direct patch to the FSW RAM should only change a variable, not an instruction or a constant. Patching an instruction or a constant changes the FSW checksum. This is important because were a flight computer reset to occur, this checksum must match the original FSW image checksum, otherwise the AFC would not progress from ROM to RAM. In February of 2003, this did happen once in-flight (to the Backup AFC) during a planned FSW upload. The problem was traced to the patching of a constant prior to 2003 that altered the FSW checksum. After this event, a flight rule was created to always use the AFC and SSR (solid state recorder) "patch table" (limited to 304 words) if a constant or instruction needs to be changed. The patch table method is robust and ensures a consistent checksum, a new FSW version number, and smooth progression into normal RAM operation, as well as preservation of the patched parameters after a flight computer reset.

In the years since arrival at Saturn, patch table use has been limited to:

1. Increase the ΔV telemetry resolution,

2. Reduce the commanded detumble acceleration to be consistent with lower RCS thrust force,
3. Adjust IRU scale factor errors based on in-flight gyro calibrations,^{62–63}
4. Change the secondary pointing vector (SRU orientation) of the safing attitude,
5. Adjust a fault protection parameter to reduce the risk of a thruster branch swap during low-altitude Titan flybys, and
6. Reduce the default spacecraft mass and RCS thrust magnitude to be consistent with the latter stages of the mission at Saturn.

A goal of the design architecture of command processing should be to minimize in-flight patching of the flight software. Parameters should be changeable through standard goal-oriented commands. On Cassini, flight software loading, latch valve states, mass properties, thrust characteristics, turn angles, power states, equipment prime-ness, burn execution, vector updates, health states, deadband changes, RWA momentum changes, channelized telemetry sample rate modifications, all utilize standard commands with full parsing, handshake, and validity checks.

Automating as many checks and tests as possible are vital to mission success. No matter how many checks, though, it remains true that engineering judgment – having experienced, well-trained engineers who have as much knowledge at their fingertips as possible – is always vital. An engineer should not just “turn the crank” like a drone. Nothing breeds success like thoughtful and engaged engineers, open to new ideas, doing their job carefully and with adequate time to do a good job.

SUMMARY

A sophisticated interplanetary spacecraft, Cassini/Huygens was launched on October 15, 1997, and arrived at Saturn on June 30, 2004. Over the past two decades, the flight performance of the Cassini attitude and articulation control subsystem has been superb. All key mission and science pointing accuracy requirements are met with significant margins. Details of the Cassini AACS designs and flight performance have been documented in more than 97 papers published in 1992–2018. To date, Cassini has returned more than 599 gigabytes of science data and 379,300 images. Together, more than 3,616 science papers have been published by scientists from 27 participating nations. Over the next several months, Cassini will send back new science data from 22 proximal orbits before the mission ends on September 15, 2017. Refs. 66–71 contain these “end of mission” data.

AACS-centric lessons learned from two decades of mission operations are documented in this paper. In retrospect, the foundations of the accomplishments of the AACS operation team come from the use of many outstanding technical innovations in the Cassini AACS design. Examples include the use of the IVP pointing model, the inclusion of a RWA drag torque estimator in the RWA controller design, the use of constraint monitor to guard against erroneous pointing commands that missed ground detection, and others. Post-launch, the mission operation team developed and used equally innovative ground software tools (such as IVP, RBOT, MAS, SEQGEN, and KPT) to automate all key processes and greatly improve the team productivity. To guard against “pilot” errors, the AACS team stressed the importance of “check, double check, and check again.” This is accomplished via the uses of flight rules, pre-uplink command testing, peer reviews of uplinkable products, completion of checklist, staffing the development of each command sequence with both prime and backup engineers, and others. Lessons learned from the Cassini operation, as well as those reported by other missions (such as the Mars Reconnaissance Orbiter⁶⁵) should be studied by mission operation teams of future spacecraft missions.

ACKNOWLEDGMENTS

The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. We wish to thank our colleagues David M. Bates and Todd Brown for their reviews of an earlier version of this paper. Any remaining errors of fact or interpretation are of course the responsibility of the authors.

REFERENCES

- ¹ Lee, A.Y. and Hanover, G., “Cassini Spacecraft Attitude Control System Flight Performance,” AIAA-2005-6269, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ² Burk, T.A., “Cassini at Saturn Proximal Orbits – Attitude Control Challenge,” AIAA-2013-4710, AIAA Guidance, Navigation, and Control Conference, Boston, Massachusetts, August 19–22, 2013.
- ³ Thomas, V.C., Alexander, J.A., Dennison, E.W., Waddell, P.G., Borghi, G., and Procopio, D., “Cassini Star Tracking and Identification Architecture,” SPIE, Volume 2221, July 5, 1994, pp. 288–298.
- ⁴ Alexander, J.W. and Chang, D.H., “Cassini Star Tracking and Identification Algorithms, Scene Simulation, and Testing,” SPIE, Volume 2803, 1996, pp. 311–336.
- ⁵ Lee, A.Y., “Risk Assessment of Cassini Sun Sensor Performance Degradation due to Hypervelocity Impact of Ring Dust Particles,” Paper AIAA 2016-2086, AIAA SciTech, January 4–8, 2016, San Diego, California.
- ⁶ Litty, E.C., Gresham, L.L., Toole, P.A., and Beisecker, D.A., “Hemispherical Resonator Gyroscope: An IRU for Cassini,” SPIE Volume 2803, October 7, 1996, pp. 299–310.
- ⁷ Barber, T. J. and Cowley, R.T., “Initial Cassini Propulsion System Inflight Characterization,” 38th Joint Propulsion Conference and Exhibit, Indianapolis, IN, 2002.
- ⁸ Bates, D. and Lee, A.Y., “In-flight Characterization of the Cassini Spacecraft Attitude Control Thrusters,” AIAA-2007-6342, AIAA Guidance, Navigation, and Control Conference, August 20–23, 2007, Hilton Head, South Carolina.
- ⁹ Burk, T.A., “Extended Bright Bodies – Flight and Ground Software Challenges on the Cassini Mission at Saturn,” Paper AIAA 2016-2090, AIAA SciTech, January 4–8, 2016, San Diego, California.
- ¹⁰ Wong, E. and Breckenridge, W., “An Attitude Control Design for the Cassini Spacecraft,” Proceedings of the AIAA Guidance, Navigation, and Control Conference, AIAA, Washington, DC, 1995, pp. 931-945.
- ¹¹ Macala, G.A., “Design of the Reaction Wheel Attitude Control System for the Cassini Spacecraft,” Paper AAS 02-121, AAS/AIAA Space Flight Mechanics Meeting, San Antonio, Texas, January 27–30, 2002.
- ¹² Lee, C. and Lee, A.Y., “Cassini Reaction Wheel Momentum Bias Optimization Tool,” AIAA-2005-6271, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ¹³ Enright, P.J., “Attitude Control of the Cassini Spacecraft During Propulsive Maneuvers,” Paper AAS 93-552, AAS/AIAA Astrodynamics Specialist Conference, Victoria, B. C., Canada, August 1993.
- ¹⁴ Burk, T.A., “Attitude Control Performance During Cassini Trajectory Correction Maneuvers,” AIAA-2005-6270, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ¹⁵ Burk, T.A., “Cassini Orbit Trim Maneuvers at Saturn – Overview of Attitude Control Flight Operations” Paper AIAA-2011-6549, AIAA Guidance, Navigation, and Control Conference, August 8–11, 2011, Portland, Oregon, 2011.
- ¹⁶ Lam, D.C., Friberg, K.H., Brown, J.M., Sarani, S., and Lee, A.Y., “An Energy Burn Algorithm for Cassini Saturn Orbit Insertion,” AIAA-2005-5994, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ¹⁷ Enright, P.J., “Thrust Vector Control Algorithm Design for the Cassini Spacecraft,” Paper AIAA 93-1043, AIAA/AHS/ASEE Aerospace Design Conference, Irvine, California, February 1993.
- ¹⁸ Rasmussen, R.D., Singh, G., Rathbun, D.B., and Macala, G.A., “Behavioral Model Pointing on Cassini Using Target Vectors,” Annual Rocky Mountain Guidance and Control Conference, Keystone, Colorado, February 1–5, 1995, pp. 91–110. See also SPIE Vol. 2803, pp. 271-287.
- ¹⁹ Burk, T.A., “A Cassini Pointing Operations Flight Experience Using Inertial Vector Propagation,” AIAA-2007-6339, AIAA Guidance, Navigation, and Control Conference, August 20–23, 2007, Hilton Head, South Carolina.
- ²⁰ Singh, G., Macala, G.A., Wong, E.C., and Rasmussen, R.D., “A Constraint Monitor Algorithm for the Cassini Spacecraft,” AIAA Paper 97-3526, AIAA Guidance, Navigation, and Control Conference, New Orleans, Louisiana, August 11–13, 1997.

- ²¹ Burk, T.A. and Bates, D.M., "Cassini Attitude Control Operation: Flight Rules and How They are enforced," AIAA-2008-6808, AIAA Guidance, Navigation, and Control Conference, August 18–21, 2008, Honolulu, Hawaii.
- ²² Brown, J., "Cassini Attitude Control Flight Software: From Development to In-flight Operation," AIAA-2008-6806, AIAA Guidance, Navigation, and Control Conference, August 18–21, 2008, Honolulu, Hawaii.
- ²³ Pilinski, E. and Lee, A.Y., "Pointing Stability Performance of the Cassini Spacecraft," *Journal of Spacecraft and Rockets*, Volume 46, No. 5, September–October, 2009, pp. 1007–1015.
- ²⁴ Lee, A.Y. and Wang, E.K., "Inflight Performance of Cassini Reaction Wheel Bearing Drag in 1997–2013," *Journal of Spacecraft and Rockets*, Vol. 52, No. 2, March–April, 2015, pp. 470–480.
- ²⁵ Brown, T., "Inflight Tuning of the Cassini RCS Attitude Controller," Paper AIAA-2011-6550, AIAA Guidance, Navigation, and Control Conference, August 8–11, 2011, Portland, Oregon, 2011.
- ²⁶ Feldman, A., Brown, J.M., Wang, E.K., Peer, S.G., and Lee, A.Y., "Reconstruction of Titan Atmosphere Density Using Cassini Attitude Control Flight Data," AAS 07-187, AAS/AIAA Space Flight Mechanics Meeting, January 28–February 1, 2007, Sedona, Arizona.
- ²⁷ Sarani, S., "Titan Atmospheric Density Reconstruction Using Cassini Guidance, Navigation, and Control Data," AIAA-2009-5763, AIAA Guidance, Navigation, and Control Conference, August 10–13, 2009, Chicago, Illinois.
- ²⁸ Lee, A.Y. and Lim, Ryan S., "Evidence of Temporal Variation of Titan Atmospheric Density in 2005–2013," AIAA-2013-4709, AIAA Guidance, Navigation, and Control Conference, Boston, Massachusetts, August 19–22, 2013.
- ²⁹ Andrade, L.C., Burk, T.A., and Pelletier, F., "Titan Density Reconstruction Using Radiometric and Cassini Attitude Control Flight Data," AIAA SciTech, January 5–9, 2015, Kissimmee, Florida, Paper AIAA 2015-0079.
- ³⁰ Teolis, B.D., Niemann, H.B., Waite, J.H., Gell, D.A., Perryman, R.S., Kasprzak, W.T., Mandt K.E., Yelle, R.V., Lee, A.Y., Pelletier, F.J., Miller, G.P., Young, D.T., Bell, J.M., Magee, B.A., Patrick, E.L., Grimes, J., Fletcher, G.G., Vuitton, V., "A Revised Sensitivity Model for Cassini INMS: Results at Titan," *Space Science Reviews*, Volume 190, Issue 1, July 2015, pp. 47–84.
- ³¹ Sarani, S., "Enceladus Plume Density Modeling and Reconstruction for Cassini Attitude Control System," AIAA 2010-2035, Space Operations 2010, April 26–30, 2010, Marshall Space Flight Center, Huntsville, Alabama.
- ³² Lee, A.Y., Wang, E.K., Pilinski, E.B., Macala, G.A., and Feldman, A., "Estimation and Modeling of Enceladus Plume Jet Density Using Cassini Flight Data," *Journal of Spacecraft and Rockets*, Vol. 50, No. 2, March–April, 2013.
- ³³ Wang, E.K. and Lee, A.Y., "Estimation of Enceladus Plume Density Using Cassini Flight Data." Paper AIAA-2011-6551, AIAA Guidance, Navigation, and Control Conference, August 8–11, 2011, Portland, Oregon, 2011.
- ³⁴ Feldman, A. and Lee, A.Y., "In-flight Estimation of Cassini Spacecraft Inertia Tensor and Thruster Magnitude," AAS 06-102, AAS/AIAA Space Flight Mechanics Meeting, January 22–26, 2006, Tampa, Florida.
- ³⁵ Stupik, J. and Burk, T.A., "Thruster-Specific Force Estimation and Trending of Cassini Hydrazine Thrusters at Saturn," Paper AIAA 2016-2087, AIAA SciTech, January 4–8, 2016, San Diego, California.
- ³⁶ Brown, J.M., Lam, D., Burk, T.A., and Wette, M., "The Role of the Flight Software Development System Simulator throughout the Cassini Mission," AIAA-2005-6389, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ³⁷ Wang, E.K., and Brown, J., "Cassini Test Methodology for Flight Software Verification During Mission Operations," AIAA-2007-6341, AIAA Guidance, Navigation, and Control Conference, 20–23 August 2007, Hilton Head, South Carolina.
- ³⁸ Brown, J., Wang, E.K., Hernandez, J., and Lee, A.Y., "Importance of Model Simulations in Cassini In-flight Mission Events," AIAA-2009-5762, AIAA Guidance, Navigation, and Control Conference, August 10–13, 2009, Chicago, Illinois.
- ³⁹ Brown, G.M., Bernard, D.E., and Rasmussen, R.D., "Attitude and Articulation Control for the Cassini Spacecraft: A Fault Tolerance Overview," 14th AIAA/IEEE Digital Avionics System Conference, Cambridge, Massachusetts, November 5–9, 1995.

- ⁴⁰ Brown, G.M. and Johnson, S.A., "An Overview of the Fault Protection Design for the Attitude Control Subsystem of the Cassini Spacecraft," American Control Conference, June 24–26, 1998, Philadelphia, Pennsylvania.
- ⁴¹ Meakin, P.C., "Cassini Attitude Control Fault Protection: Launch to End of Prime Mission Performance," AIAA-2008-6809, AIAA Guidance, Navigation, and Control Conference, August 18–21, 2008, Honolulu, Hawaii.
- ⁴² Lee, A.Y., "Model-based Thruster Leakage Monitor for the Cassini Spacecraft," *Journal of Spacecraft and Rockets*, Vol. 36, No. 5, pp. 745–749, September–October, 1999.
- ⁴³ Cooney, L.A., "Cassini Attitude Control Fault Protection: Flight Operations Strategy Changes for Extended Mission," AIAA-2010-7561, AIAA Guidance, Navigation, and Control Conference, August 2–5, 2010, Toronto, Ontario, Canada.
- ⁴⁴ Bates, D.M., "Cassini Attitude and Articulation Control Subsystem Fault Protection Challenges during Saturn Proximal Orbits," AIAA SciTech, 5–9 January 2015, Kissimmee, Florida, Paper AIAA 2015-0077.
- ⁴⁵ Burk, T.A., "Managing Cassini Safe Mode Attitude at Saturn," AIAA-2010-7558, AIAA Guidance, Navigation, and Control Conference, August 2–5, 2010, Toronto, Ontario, Canada.
- ⁴⁶ Allestad, D.L., Standley, S.P., Chang, L., and Bone, B.D., "Systems Overview of the Cassini-Huygens Probe Relay Critical Sequence," AIAA-2005-6388, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ⁴⁷ Battin, R.H., *An Introduction to the Mathematics and Methods of Astrodynamics*, AIAA Education Series, 1987.
- ⁴⁸ Goodyear, W.H., "Completely General Closed-Form Solution for Coordinates and Partial Derivatives of the Two-Body Problem," *The Astronomical Journal*, Vol. 70, No. 3, April 1965, pp. 189–192.
- ⁴⁹ Barrodale, I. and Phillips, C., "Algorithm 495. Solution of an Overdetermined System of Linear Equations in the Chebyshev Norm," *ACM Transactions on Mathematical Software*, Vol. 1, Issue 3, September 1975, pp. 264–270.
- ⁵⁰ Velarde Yang, G., Kirby, C., and Mohr, D., "Cassini's Maneuver Automation Software (MAS) Process: How to Successfully Command 200 Navigation Maneuvers," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Honolulu, Hawaii, August 18–21, 2008.
- ⁵¹ Mizukami, M., Barber, T.J., Christodoulou, L.N., Guernsey, G.S., and Haney, W.A., "Cassini Spacecraft Reaction Control System Thrusters Flight Experience," JANNAP Propulsion Meeting, Colorado Springs, CO, May 2010.
- ⁵² Burk, T., "Avoiding Human Error in Mission Operations – Cassini Flight Experience," AIAA-2012-4607, AIAA Guidance, Navigation, and Control Conference, August 13–16, Minneapolis, Minnesota, 2012.
- ⁵³ Badaruddin, K., Hernandez, J., and Brown, J., "The Importance of Hardware-In-The-Loop Testing to the Cassini Mission to Saturn," *IEEE Aerospace Conference*, Track 12.0503, Paper 1231, Big Sky, MT, March 3–10, 2007.
- ⁵⁴ Sung, T. and Burk, T., "Extended Bright Bodies – Flight and Ground Software Challenges on the Cassini Mission at Saturn," AIAA-2016-2090, AIAA Guidance, Navigation, and Control Conference, SciTech 2016, January 4–8, 2016, San Diego, CA.
- ⁵⁵ Brown, T., "Y-Biasing: A New Operational Method for Cassini to Control Thruster Usage while Managing Reaction Wheel Momentum," AIAA-2010-7560, AIAA Guidance, Navigation, and Control Conference, August 2–5, 2010, Toronto, Ontario.
- ⁵⁶ Perkins, D., *NASA Internal Review Board Report on Mars Global Surveyor Spacecraft Loss of Contact*, April 13, 2007.
- ⁵⁷ Macala, G.A., Lee, A.Y., and Wang, E.K., "Feasibility Study of Two Candidate Reaction Wheel/Thruster Hybrid Control Architecture Designs for the Cassini Spacecraft," *Journal of Spacecraft and Rockets*, Vol. 51, No. 2, March–April, 2014, pp. 574–585.
- ⁵⁸ Bates, D.M., "Cassini Spacecraft In-flight Swap To Backup Attitude Control Thrusters," AIAA-2010-7559, AIAA Guidance, Navigation, and Control Conference, 2–5 August 2010, Toronto, Ontario, Canada.
- ⁵⁹ Cervantes, D., Badaruddin, K., and Huh, S.M., "Integrated Testing of the Saturn Orbit Insertion Critical Sequence," AIAA-2005-6272, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.

- ⁶⁰ Montañez, L., Bone, B.D., and Laufer, P., “Evolution of the Cassini Attitude and Articulation Control Subsystem Simulation During a Seven-year Cruise Mission Phase,” AIAA-2005-6273, AIAA Guidance, Navigation, and Control Conference, August 15–18, 2005, San Francisco, California.
- ⁶¹ Wagner, S., “Maneuver Performance Assessment of the Cassini Spacecraft Through Execution-Error Modeling and Analysis,” Paper AAS 14-390, AAS/AIAA Space Flight Mechanics Meeting, January 26–30, 2014, Santa Fe, New Mexico.
- ⁶² Burrough, E.L. and Lee, A.Y., “In-flight Characterization of the Cassini Spacecraft’s Inertial Reference Units,” AIAA-2007-6340, AIAA Guidance, Navigation, and Control Conference, August 20–23, 2007, Hilton Head, South Carolina.
- ⁶³ Brown, T., “Inflight Performance of the Cassini Hemispherical Quartz Resonator Gyro Inertial Reference Units,” AIAA-2013-4630, AIAA Guidance, Navigation, and Control Conference, Boston, Massachusetts, August 19–22, 2013.
- ⁶⁴ Lee, A.Y. and Stupik, J., “Inflight Characterization of the Cassini Spacecraft Propellant Slosh,” *Journal of Spacecraft and Rockets*, Volume 54, 2017. See also AIAA-2015-0078.
- ⁶⁵ Chapel, J.D., Schmitz, E., Sidney, W.P., Johnson, M.A., Good, P.G., Wynn, J.A., and Bayer, T., “Attitude Control Performance for MRO Aerobraking and Initial Science Phase,” AAS 07-097, Proceedings of the AAS Rocky Mountain Guidance and Control Conference, February 3–7, 2007, Breckenridge, Colorado.
- ⁶⁶ Lee, A.Y., “Cassini Reaction Wheel Controller Performance During Low-rpm Operations and Alternative Drag Torque Compensation Schemes,” AIAA SciTech, 8–12 January 2018, Kissimmee, Florida, AIAA Guidance, Navigation, and Control Conference.
- ⁶⁷ Sung, T.S., “Attitude Control Subsystem Performance of the Cassini Spacecraft During Proximal Ring Plane Crossings Between Saturn and the Inner D-Ring,” AIAA SciTech, 8–12 January 2018, Kissimmee, Florida, AIAA Guidance, Navigation, and Control Conference.
- ⁶⁸ Brown, T.S., “A Full Mission Summary (1997-2017) of the Attitude Control Performance during Orbit Trim Maneuvers performed by the Cassini-Huygens Spacecraft,” AIAA SciTech, 8–12 January 2018, Kissimmee, Florida, AIAA Guidance, Navigation, and Control Conference.
- ⁶⁹ Stupik, J., “Mission Summary of Cassini Spacecraft Guidance and Control Hardware Health and Performance,” AIAA SciTech, 8–12 January 2018, Kissimmee, Florida, AIAA Guidance, Navigation, and Control Conference.
- ⁷⁰ Andrade, L.G., “Skimming through Saturn’s Atmosphere: The Climax of the Cassini Grand Finale Mission,” AIAA SciTech, 8–12 January 2018, Kissimmee, Florida, AIAA Guidance, Navigation, and Control Conference.
- ⁷¹ Burk, T.A., “Enhanced Reconstructed C-Kernels from Cassini Flight Telemetry at Saturn,” AIAA SciTech, 8–12 January 2018, Kissimmee, Florida, AIAA Guidance, Navigation, and Control Conference.
- ⁷² Bruno, D., “Contingency Mixed Actuator Controller Implementation for the Dawn Asteroid Rendezvous Spacecraft,” Paper AIAA-2012-5289, AIAA Space 2012 Conference & Exposition, September 11–13, 2012, Pasadena, California.
- ⁷³ Smith, B.A., Lim, R.S., Feldman, A., and Salami, M., “Dawn Spacecraft Performance at Ceres: Results of Hybrid Controller for Ceres Mapping,” Paper AAS 17-151, AAS Guidance and Control Conference, Breckenridge, Colorado, February 2–8, 2017.

APPENDIX: AACS-CENTRIC INCIDENT-SURPRISE-ANOMALY REPORTS (ISA), 1997–2017

Cassini documents non-conformances via Incident-Surprise-Anomaly (ISA) reports. Verification, validation and corrective actions are thoroughly documented prior to anomaly report evaluation and closure by Mission Assurance. Beside ISA's that addressed AACS hardware and software problems, reports that addressed AACS ground software, operator errors, etc. are also documented. In the following table, AACS-centric ISA are grouped by categories based on the root causes of the anomalies.

ISA categories	Percent [%]
Anomalous RWA bearing drag torque	7.9
Inaccurate prediction of residual ΔV 's	5.3
Inaccurate AACS parameter values	2.0
Selections of FP and CMT parameters	4.6
AACS FSW idiosyncrasies	5.9
AACS controller performance limitations	3.9
Impacts of extended objects on tracker performance	3.9
Operator's errors	13.2
KPT ground software errors (FR violation checks, interrupted turns, RWA biasing, IVP, SID, etc.)	25.0
Procedural errors	3.3
MAS data system errors	1.3
Errors of command/telemetry dictionaries	1.3
Impacts of CDS problems on AACS performance	1.3
IVP idiosyncrasies	8.6
Engine and thruster hardware problems	0.7
Live IVP update errors	7.2
Radiation impacts on IRU performance	3.9
Unknown root cause(s)	0.7
All categories	100.0